Service Overlay Forest Embedding for Software-Defined Cloud Networks

Jian-Jhih Kuo*, Shan-Hsiang Shen*†, Ming-Hong Yang*‡, De-Nian Yang*, Ming-Jer Tsai§ and Wen-Tsuen Chen*§
*Inst. of Information Science, Academia Sinica, Taipei, Taiwan
†Dept. of Computer Science & Information Engineering,
National Taiwan University of Science & Technology, Taipei, Taiwan
‡Dept. of Computer Science & Engineering, University of Minnesota, Minneapolis MN, USA
§Dept. of Computer Science, National Tsing Hua University, Hsinchu, Taiwan
E-mail: {lajacky,sshen3,curtisyang,dnyang,chenwt}@iis.sinica.edu.tw and mjtsai@cs.nthu.edu.tw

arXiv:1703.09025v1 [cs.NI] 27 Mar 2017

***Abstract*—Network Function Virtualization (NFV) on Software-Defined Networks (SDN) can effectively optimize the allocation of Virtual Network Functions (VNFs) and the routing of network flows simultaneously. Nevertheless, most previous studies on NFV focus on unicast service chains and thereby are not scalable to support a large number of destinations in multicast. On the other hand, the allocation of VNFs has not been supported in the current SDN multicast routing algorithms. In this paper, therefore, we make the first attempt to tackle a new challenging problem for finding a service forest with multiple service trees, where each tree contains multiple VNFs required by each destination. Specifically, we formulate a new optimization, named Service Overlay Forest (SOF), to minimize the total cost of all allocated VNFs and all multicast trees in the forest. We design a new $3\rho_{ST}$-approximation algorithm to solve the problem, where ρ_{ST} denotes the best approximation ratio of the Steiner Tree problem, and the distributed implementation of the algorithm is also presented. Simulation results on real networks for data centers manifest that the proposed algorithm outperforms the existing ones by over 25%. Moreover, the implementation of an experimental SDN with HP OpenFlow switches indicates that SOF can significantly improve the QoE of the Youtube service.**

I. Introduction

The media industry is now experiencing a major change that alters user subscription patterns and thereby inspires the architects to rethink the design [1]. For example, the live video streaming on Anvato [2] enables online video editing for content providers, ad insertion for advertisers, caching, and transcoding for heterogeneous user devices. Google has acquired Anvato with the above abundant functions and integrated its architecture into Google Cloud to develop the next-generation Youtube.[i] Therefore, it is envisaged that the next-generation Youtube requires more computation functionalities and resources in the cloud. For distributed collaborative virtual reality (VR), it is also crucial to allocate distributed computation resources for important tasks such as collision detection, geometric constraint matching, synchronization, view consistency, concurrency and interest management [3]–[5]. Network Function Virtualization (NFV) has been regarded as a promising way [1], [6] that exploits Virtual Machines (VMs) to divide the required function into building blocks connected with a service chain [7]. A service chain passes through a set of Virtual Network Functions (VNFs) in sequence, and Netflix [8] has adopted AWS [9] to support the service chains. Current commercial solutions usually assign an individual service chain for each end user for unicast [10]–[13]. Nevertheless, it is expected that this approach is not scalable because duplicated VNFs and network traffic are involved to serve all users if they require the same content, such as live/linear content broadcast. The global consumer research [14] manifests that although the unicast video on demand becomes more and more popular, the live/linear content broadcast and multicast nowadays still account for over 50% of viewing hours per week, from companies such as Sony Crackle [15] and Pluto TV [16], because it effectively attracts the users through a shared social experience to instantly access the contents. However, currently there is no effective solution to support large-scale content distributions with abundant computation functionalities for content providers and end users.

For scalable one-to-many communications, multicast exploits a tree to replicate the packets in branching routers. Compared with unicast flows, a multicast tree can effectively reduce the bandwidth consumption in backbone networks by over 50% [17], especially for multimedia traffic [18]. Currently, shortest-path trees are employed by Internet standards (such as PIM-SM [19]) because they can be efficiently constructed in a distributed manner. Nevertheless, the routing is not flexible since the path from the source to each destination needs to follow the corresponding shortest path. Recently, the flexible routing for traffic engineering becomes increasingly important with the emergence of Software-Defined Networks (SDNs), whereas centralized computation can be facilitated in an SDN controller to find the optimal routing, such as Steiner Tree [20] in Graph Theory or its variations [21], [22]. Thus, multicast traffic engineering has been regarded very promising for SDNs.

Nevertheless, the above approaches and other existing multicast routing algorithms [23], [24] are not designed to support NFV because the nodes (e.g., the source and destinations) that need to be connected in a tree are specified as the problem input. On the contrary, here VNFs are also required to be spanned in a tree for NFV, and the problem is more challenging since VMs also need to be selected, instead of being assigned as the problem input. Moreover, multicast NFV indeed is more complicated when it is necessary to employ multiple multicast trees as a forest for a group of destinations, and this feature is crucial for Content Deliver Networks (CDNs) with multiple video source servers. In this case, the video source also needs to be chosen for each end user [25].

In this paper, therefore, we make the first attempt to explore the resource allocation problem (i.e., both the VM selection,

[i]http://www.msn.com/en-us/news/technology/google-buys-a-backbone-for-pay-tv-services/ar-BBu61eB?li=AA4Zoy&ocid=spartanntp

source selection, and the tree routing) for a *service forest* involving multiple multicast trees, where the path from a source to each destination needs to traverse a sequence of demanded services (i.e., a service chain) in the tree.[ii] We formulate a new optimization problem for multi-tree NFV in software defined cloud networks, named, Service Overlay Forest (SOF), to minimize the total cost of the selected VMs and trees. Given the sources and the destinations with each destination requiring a chain of services, the SOF problem aims at finding an overlay forest that 1) connects each destination to a source and 2) visits the demanded services in selected VMs in sequence before reaching the destinations.

Fig. 1 first compares a service tree and a service forest. Fig. 1(a) is the input network with the cost labeled beside each node and edge to represent the link connection cost and the VM setup cost, respectively. Assume that there are two destinations 9 and 10, and their demanded service chain consists of two VNFs, f_1 and f_2 in order. A Steiner tree in Fig. 1(b) spanning source node 1 and both destinations incurs the total cost as 34 if VMs 2 and 3 are employed. Note that the edge between VMs 2 and 3 is visited twice to reach destination 10, and the cost of the edge is thus required to be included twice. More specifically, the edge costs from source 1 to VM 3 (f_1), from VM 3 to VM 2 (f_2), and from VM 2 to destinations 9, 10, are 1, 3, $20+1+1+1+3+1+1=28$, respectively. Thus, the edge cost is $1+3+28=32$ while the node cost is $1+1=2$. By contrast, the cost of a service forest with two trees and four VMs selected is 14 in Fig. 1(c), which significantly reduces the cost by about 60 %. This example manifests that consolidating the services in few VMs may not always lead to the smallest cost because the edges to connect multiple destinations are also important. Therefore, multiple trees with multiple sources are promising to further reduce the cost.[iii]

In this paper, we first prove that the problem is NP-hard. To investigate the problem in depth, we will step-by-step reveal the thinking process of the algorithm design from the single-source case to the general case, and then propose a $3\rho_{ST}$-approximation algorithm,[iv] named Service Overlay Forest Deployment Algorithm (SOFDA) for the general case, where ρ_{ST} denotes the best approximation ratio of the Steiner Tree problem (e.g., the current best one is 1.39). The single-source case is more difficult than the traditional Steiner tree problem because not only the terminal nodes (i.e., source and destinations) need to be spanned but also a set of VMs is required to be selected and spanned to install VNFs in sequence. Also, the general case is more challenging than the single-source case, because a service tree is necessary to be created for each source, and

[ii] In this paper, we first consider the static multicast, and then clarify the static case is already a good step forward and discuss how to adapt the proposed algorithm to the dynamic case in Sections VII-A and VII-C, respectively.

[iii] In this paper, we assume that the setup cost for a source node is negligible. The source with the setup cost is further discussed in Appendix D.

[iv] Compared with the traditional Steiner Tree problem, the problem considered in this paper is more difficult due to new SDN/NFV constraints involved. Indeed, several recent research works [21], [22] on SDN multicast and NFV service chain embedding (e.g., [11], [13]) have massive approximation ratios (e.g. $O(|D|)$, where $|D|$ denotes the number of destinations, and $O(|\mathcal{C}|)$, where $|\mathcal{C}|$ denotes the length of demanded service chain). By contrast, the approximation ratio of this paper is $3\rho_{ST}$, where ρ_{ST} is the best approximation ratio of the Steiner Tree problem (e.g., the current best one is 1.39), which is smaller than the above works. Moreover, the simulation results manifest that empirically the performance is very close to the optimal solutions obtained by the proposed Integer Programming formulation.

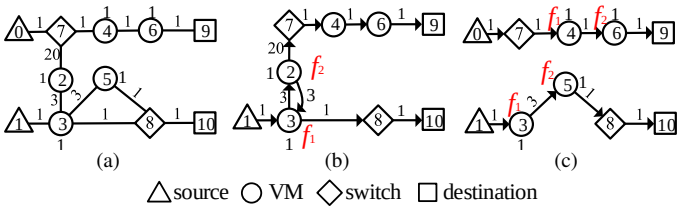

Fig. 1. Comparison of service trees and service forests. (a) Input network. (b) Steiner tree with predetermined VMs. (c) Service Overlay Forest.

the *VNF conflict* (i.e., a VM is allocated with too many VNFs from multiple trees) tends to occur in this case.

Therefore, SOFDA is designed to 1) assign multiple sources for varied trees with multiple VMs, 2) allocate the VMs for each tree to provide a service chain for each destination, and 3) find the routing of each tree to span the selected source, VMs, and destinations. Simulation on real topologies manifests that SOFDA can effectively reduce the total cost for data center networks. In addition, a distributed SOFDA is proposed to support the multi-controller SDNs. Implementation on an experimental SDN for Youtube traffic also indicates that the user QoE can be significantly improved for transcoded and watermarked video streams. The rest of this paper is organized as follows. The related works are summarized in Section II. We formulate the SOF problem in Section III and design the approximation algorithms in Sections IV and V. The distributed algorithm is presented in Section VI. Some important issues are discussed in Section VII. The simulation and implementation results are presented in Section VIII, and we conclude this paper in Section IX.

II. Related Work

Traffic engineering for unicast service chains in SDN has drawn increasing attention recently. Lukovszki *et al.* [11] point out that the length of a service chain is necessary to be bounded and present an efficient online algorithm to maximize the number of deployed service chains, whereas the maximal number of VMs hosted on a node is also guaranteed. Xia *et al.* [12] jointly consider the optical and electrical domains and minimize the number of domain conversions in all service chains. Moreover, Kuo *et al.* [13] strike the balance between link utilization and server usage to maximize the total benefit. Nevertheless, the above studies only explore unicast routing for service chains and do not support multicast.

Multicast traffic engineering for SDN is more complicated than traditional unicast traffic engineering. Huang *et al.* [26] first incorporate the flow table scalability in the design of the multicast tree routing in SDN. Shen *et al.* [21] then further consider the packet loss recovery in reliable multicast routing for SDN. Recently, the routing of multiple trees in SDN [22] has been studied to ensure that the routing follows both the link capacity and the TCAM size. The problem is more challenging due to the above two constraints, and the best approximation ratio that can be achieved is only D (i.e., the maximum number of destinations in a tree). However, the dimension of service allocation in VMs has not been explored in the above work. Recently, special cases on a tree [27], [28] with only one source and one VM have been explored. Overall, the above approaches are not designed to support a service forest with multiple VNFs and multiple trees, and the problem here is more challenging because *VNF conflict* due to the overlapping of trees will occur.

To the best knowledge of the authors, this paper is the first one that explores both routing and VM selection for multiple trees to construct a forest in SDN. As explained in Section I, the service forest is important for many emerging and crucial multimedia applications in CDN that require intensive cloud computing.

III. The Service Overlay Forest Problem

A service overlay forest consists of a set of service overlay trees. Each service overlay tree spans one source, a set of VMs for enabled VNFs, and a subset of destinations. Any two service overlay trees do not overlap since each destination only needs to connect to a source via a service chain in a tree. In the following, we first formally define the problem. We are given:

1) a network $G = \{V = M \cup U, E\}$, where each link $e \in E$ is associated with a nonnegative cost $c(e)$ denoting the connection cost of link e to forward the demand of destinations, each virtual machine (VM) $v \in M$ is associated with a nonnegative cost $c(v)$ denoting the setup cost of VM v to run a virtual network function (VNF), and each switch $v \in U$ is associated with cost 0,
2) a set of destinations $D \subseteq V$ requesting the same demand,
3) a set of sources $S \subseteq V$ having the demands of destinations, and
4) a chain of VNFs $\mathcal{C} = (f_1, f_2, \cdots, f_{|\mathcal{C}|})$ required to process the demand of destinations.

The **S**ervice **O**verlay **F**orest (SOF) problem is to construct a service overlay forest consisting of the service overlay trees with the roots in S, the leaves in D, and the remaining nodes in V, so that there exists a chain of VNFs from a source to each destination. A chain of VNFs is represented by a *walk*, which is allowed to traverse a node (i.e., a VM or a switch) multiple times. In each walk, a *clone* of a node and the corresponding incident links are created to foster an additional one-time pass of the node, and only one of its clones is allowed to run VNF to avoid duplicated counting of the setup cost. For example, in the second feasible forest (colored with light gray) of Fig. 2(b), a walk from source 1 to destination 8 passes VM 2 twice without running any VNF, and there are two clones of VM 2 on the walk. For each destination $t \in D$, SOF needs to ensure that there exists a path with clone nodes (i.e., a walk on the original G) on which $f_1, f_2, \cdots, f_{|\mathcal{C}|}$ are performed in sequence from a source $s \in S$ to t in the service overlay forest.

The objective of SOF is to minimize the total setup and connection cost of the service overlay forest, where the setup and connection costs denote the total cost of the VMs and links, respectively. Note that the cost of a link in G is counted twice if the link is duplicated because its terminal nodes are cloned. In this paper, it is assumed that a VM can run at most one VNF in the network G. The scenario that requires a VM to support multiple VNFs can be simply addressed by first replicating the VM multiple times in the input graph G.

Example 1. Fig. 2 presents three examples for the service overlay forests. The first service overlay forest consists of two service overlay trees, where the demand of destination 8 (or 9) is routed from source 1 (or 0) along the walk $(1, 2, 4, 10, 6, 8)$ (or $0, 3, 11, 5, 7, 9$), and the demand is processed by VNFs f_1 and f_2 at VMs 4 and 6 (or 3 and 7), respectively. The total cost of the first service overlay forest is 82, where the setup cost and connection cost are 50 and 32, respectively. In

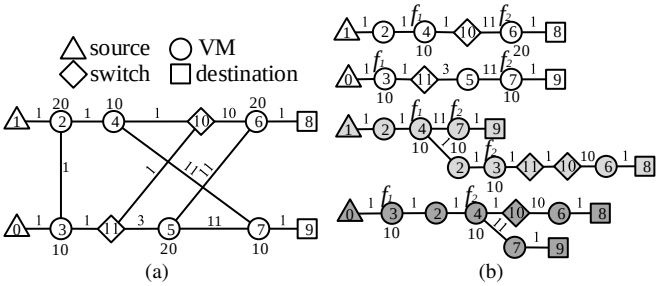

Fig. 2. Example of service overlay forests. (a) The input network G. (b) The service overlay forests with $\mathcal{C} = (f_1, f_2)$ constructed for G.

the second service overlay forest (including only one tree), source 1 first routes the demand to VM 4 for VNF f_1. Subsequently, VM 4 forwards the demand to VM 7 and VM 2 for VNF f_2, respectively. Finally, the demand is forwarded towards destinations 8 and 9, respectively. The setup cost and connection costs of the second service overlay forest are 30 and 29, respectively. In the third service overlay forest (tree), the demand is first routed from source 1 to VM 3 for VNF f_1 and then toward VM 4 for VNF f_2, and finally to destinations 8 and 9, respectively. The third service overlay forest is an optimal service overlay forest with the setup cost and connection cost as 20 and 27, respectively. ■

A. Integer Programming

In the following, we present the Integer Programming (IP) formulation for SOF. Our IP formulation first identifies the service chain for each destination and then constructs the whole service forest accordingly. To find the walk of the service chain, we first assign the VMs corresponding to each VNF in the walk and then find the routing of the walk between every two consecutive VMs. More specifically, SOF includes the following binary decision variables. Let $\gamma_{d,f,u}$ denote if node u is assigned as the enabled VM for VNF f in the walk to destination d. Let $\pi_{d,f,u,v}$ denote if edge $e_{u,v}$ is located in the walk connecting the enabled VM of VNF f and the enabled VM of the next VNF f_N. Note that the above walk will belong to a service tree rooted at the enabled candidate node of VNF f. Therefore, to find the service forest for f, let binary variable $\tau_{f,u,v}$ represent if edge $e_{u,v}$ is located in the forest. On the other hand, binary variable $\sigma_{f,u}$ represents if node u is assigned as the enabled VM of service f for the whole service forest. Notice that each destination d may desire a different VNF f on the same enabled VM u according to $\gamma_{d,f,u}$, but the constraint later in this section will ensure that only one VNF is allowed to be allocated to u by properly assigning $\sigma_{f,u}$ accordingly.

The objective function for SOF is as follows.

$$\min \sum_{f \in \mathcal{C}} \sum_{u \in V} c(u)\sigma_{f,u} + \min \sum_{f \in \mathcal{C}} \sum_{e_{u,v} \in E} c(e_{u,v})\tau_{f,u,v},$$

where the first term represents the total setup cost of all VMs, and the second term is the connection cost of the service forest. The IP formulation contains the following constraints.

1) Service Chain Constraint. The following four constraints first assign the enabled VM for each service chain.

$$\sum_{s\in S}\gamma_{d,f_S,s}=1,\ \forall d\in D, \quad (1)$$
$$\sum_{u\in M}\gamma_{d,f,u}=1,\ \forall d\in D, f\in\mathcal{C}, \quad (2)$$
$$\gamma_{d,f_D,d}=1,\ \forall d\in D, \quad (3)$$
$$\gamma_{d,f_D,u}=0,\ \forall d\in D, u\in V-\{d\}. \quad (4)$$

Constraint (1) ensures that each destination chooses one source s in S as its service source, where f_S denotes the function as the source of the service chain. Constraint (2) finds a node u from M as the enabled VM of each VNF f for each destination. Constraints (3) and (4) assign only destination d for function f_D, where f_D denotes the function as the destination of the service chain. Here notations f_S and f_D are incorporated in our IP formulation in order to support the routing constraints described later. In other words, a service chain traverses the nodes with $f_S, f_1, ..., f_{|\mathcal{C}|}, f_D$ sequentially.

2) Service Forest Constraint. The following two constraints assign the enabled VM for the whole service forest.

$$\gamma_{d,f,u}\le\sigma_{f,u},\ \forall d\in D, f\in\mathcal{C}, u\in V, \quad (5)$$
$$\sum_{f\in\mathcal{C}}\sigma_{f,u}\le 1,\ \forall u\in V, \quad (6)$$

Constraint (5) assigns u as the enabled VM of VNF f for the whole service forest if u has been selected by at least one destination d for VNF f. Constraint (6) ensures that each node u is in charge of at most one VNF.

3) Chain and Forest Routing Constraints. The following two constraints find the routing of the whole service forest.

$$\sum_{v\in N_u}\pi_{d,f,u,v}-\sum_{v\in N_u}\pi_{d,f,v,u}\ge\gamma_{d,f,u}-\gamma_{d,f_N,u},$$
$$\forall d\in D, f\in\mathcal{C}\cup\{f_S\}, u\in V, \quad (7)$$
$$\pi_{d,f,u,v}\le\tau_{f,u,v},\ \forall d\in D, f\in\mathcal{C}\cup\{f_S\}, e_{u,v}\in E. \quad (8)$$

Constraint (7) is the most complicated one. It first finds the routing of the service chain for each destination d. For the source u of a service chain, $\gamma_{d,f_S,u}=1$ and $\gamma_{d,f_N,u}=\gamma_{d,f_1,u}=0$, where f_1 is the first VNF in $\mathcal{C}$. In this case, the constraint becomes

$$\sum_{v\in N_u}\pi_{d,f_S,u,v}-\sum_{v\in N_u}\pi_{d,f_S,v,u}\ge 1.$$

It ensures that at least one edge $e_{u,v}$ incident from u is selected for the service chain because no edge $e_{v,u}$ incident to u is chosen (i.e., $\sum_{v\in N_u}\pi_{d,f_S,v,u}=0$ for the source u). By contrast, for any intermediate switch u in the walk from the source to the enabled VM of f_1, $\gamma_{d,f_S,u}=0$ and $\gamma_{d,f_1,u}=0$, and the constraint becomes

$$\sum_{v\in N_u}\pi_{d,f_S,v,u}\le\sum_{v\in N_u}\pi_{d,f_S,u,v}.$$

When any edge $e_{v,u}$ incident to u has been chosen in the walk, the above constraint states that at least one edge $e_{u,v}$ incident from u must also be selected in order to construct the service chain iteratively. The above induction starts from the source of the walk to the previous node of the enabled VM of f_1. Afterward, for the enabled VM u of f_1, $\gamma_{d,f_1,u}=1$ and $\gamma_{d,f_S,u}=0$, and the constraint becomes

$$\sum_{v\in N_u}\pi_{d,f_S,v,u}-\sum_{v\in N_u}\pi_{d,f_S,u,v}\le 1.$$

Since $\pi_{d,f_S,v,u}=1$ for only one edge $e_{v,u}$ in the walk incident to u, the above constraint is identical to $\sum_{v\in N_u}\pi_{d,f_S,u,v}\ge 0$. Therefore, $\pi_{d,f_S,u,v}$ is allowed to be 0 for every edge $e_{u,v}$ to minimize the objective function, implying that no data of f_S will be sent from the enabled VM of f_1. By contrast, $\pi_{d,f_1,u,v}$ will be 1 for one edge $e_{u,v}$ due to constraint (6), implying that the enabled VM of f_1 will deliver the data in one edge incident from u, and the above induction repeats sequentially for every service f in $\mathcal{C}$ until it reaches the destination d. Finally, constraint (8) states that any edge $e_{u,v}$ is in the service forest if it is in the service chain for at least one destination d.

B. The Hardness

The SOF problem is NP-hard since a metric version of the Steiner Tree problem (see Definition 1) can be reduced to the SF problem in polynomial time. The complete proof is presented in Appendix A.

Definition 1. [20] Given a weighted graph $G=\{V,E\}$ with edge costs, a root $r\in V$ and a node set $U\subseteq V\setminus\{r\}$, a Steiner Tree is a minimum spanning tree that roots at s and spans all the nodes in U, where $U\neq\emptyset$.

Theorem 1. *The SOF problem is NP-hard.*

IV. Special Case with Single Tree

In this subsection, we propose a $(2+\rho_{ST})$-approximation algorithm, named Service Overlay Forest Deployment Algorithm with Single Source (SOFDA-SS) to explore the fundamental characteristics of the problem, and a more complicated algorithm for the general case with multiple sources will be presented in the next section. SOFDA-SS includes the following two phases. The first phase chooses the most suitable VM to install the last VNF (i.e., called last VM in the rest of this paper) and then finds a minimum-cost service chain[v] between the source and the last VM. Afterward, the second phase finds a minimum-cost Steiner tree to span the VM and all the destinations. The selection of the last VM is crucial due to the following trade-offs. Choosing a VM closer to the source tends to generate a shorter service chain, but it may create a larger service tree if the last VM is distant from all destinations. Also, it is important to address the trade-off between the setup cost and connection cost, because a VM with a smaller setup cost will sometimes generate a larger tree. The pseudo code of SOFDA-SS is presented in Appendix E (see Algorithm 1) .

Therefore, to achieve the approximation ratio, it is necessary for SOFDA-SS to carefully examine every possible VM to derive a Steiner tree and evaluate the corresponding cost. For every VM u, to obtain a walk W_G (i.e., service chain) from source s to u with $|\mathcal{C}|$ VMs (so that the VNFs $f_1, f_2, \cdots, f_{|\mathcal{C}|}$ can be installed in sequence) in G, we first propose a graph transformation from G to $\mathcal{G}$ and then find the k-stroll [29] from s to u defined as follows.

Definition 2. Given a weighted graph $\mathcal{G}=\{\mathcal{V},\mathcal{E}\}$ and two nodes s and u in $\mathcal{V}$, the k-stroll problem is to find the shortest walk that visits at least k distinct nodes (including s and u) from s to u in $\mathcal{G}$.

[v]The next section will extend the service chain into a service tree with multiple last VMs.

SOFDA-SS constructs an instance $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of the k-stroll problem from G as follows. Let $\mathcal{V}$ consist of s and all VMs in G (i.e., $\mathcal{V} = M \cup \{s\}$). Let $\mathcal{E}$ contain all edges between any two nodes in $\mathcal{V}$ (i.e., $\mathcal{G}$ is a complete graph). The cost of the edge between nodes v_1 and v_2 in $\mathcal{E}$ is defined as follows,

$$c(v_1, v_2) = \sum_{(a,b)\in P} c((a,b)) + \begin{cases} \frac{c(u)+c(v_2)}{2} & \text{if } v_1 = s, \\ \frac{c(v_1)+c(u)}{2} & \text{else if } v_2 = s, \\ \frac{c(v_1)+c(v_2)}{2} & \text{otherwise,} \end{cases}$$

where u and P denote the last VM and the shortest path between nodes v_1 and v_2 in G, respectively. In other words, the cost of each shortest path in G is first included in the cost of the corresponding edge in $\mathcal{E}$. Afterward, since the data always enter and leave the VM running an intermediate VNF ($\neq f_{|\mathcal{C}|}$), the setup cost of the VM is shared by the incoming and outgoing edges of the VM. Finally, the setup cost of last VM u is shared by the outgoing edge of s and the incoming edge of u. The edge costs of $\mathcal{G}$ are assigned in the above way to ensure that the shortest walk with $|\mathcal{C}|$ VMs in G is identical to the shortest path with $|\mathcal{C}| + 1$ nodes in $\mathcal{G}$.

Clearly, $\mathcal{G}$ can be constructed in polynomial time. Then, SOFDA-SS finds a k-stroll walk $W'_{\mathcal{G}}$ that visits exactly $|\mathcal{C}|+1$ distinct nodes from source s to the last VM u (i.e., $k = |\mathcal{C}|+1$) in $\mathcal{G}$. Then, SOFDA-SS finds the corresponding walk W_G (i.e., a service chain from s to u in G) that visits exactly $|\mathcal{C}|$ distinct VMs in G by concatenating each shortest path corresponding to a selected edge in $|\mathcal{C}|$, and each path connects two consecutive nodes, u_j and u_{j+1}, on walk $W_{\mathcal{G}}$, where $1 \le j \le |\mathcal{C}|$. Finally, the demanded VNFs $f_1, f_2, ..., f_{|\mathcal{C}|}$ can be deployed in order on the walk with $|\mathcal{C}|$ VMs from s to u.

Example 2. Fig. 3 presents an illustrative example for SOFDA-SS. First, for VM 7, the walk $W_G = (u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$ with $u_1 = 1$ and $u_{|\mathcal{C}|+1} = 7$ is obtained as follows. An instance $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of the k-stroll problem is first constructed with $s = 1$, $M = \{2, 3, 4, 5, 6, 7\}$, and $u = 7$, where $\mathcal{V}$ is set to $\{1, 2, 3, 4, 5, 6, 7\}$, $\mathcal{E}$ is set as $\{(x, y)|x, y \in \mathcal{V}\}$, the cost of the edge between nodes 1 and 6 is set to $c((1, 2)) + c((2, 4)) + c((4, 6)) + \frac{c(5)+c(6)}{2} = 14$, and the cost of the edge between nodes 2 and 6 is set as $c((2, 4)) + c((4, 6)) + \frac{c(2)+c(6)}{2} = 13$. Subsequently, we acquire a walk $W'_{\mathcal{G}} = W_{\mathcal{G}} = (1, 2, 4, 3, 5, 7)$ in $\mathcal{G}$ and the corresponding walk $W_G = (1, 2, 4, 2, 3, 5, 7)$ in G. After W_G is obtained, the service overlay forest with the last VM (i.e., 7) is constructed, where the demand is first routed from source 1 to VM 7 along the walk $W_G = (1, 2, 4, 2, 3, 5, 7)$, and f_1, f_2, f_3,f_4, and f_5 is processed at VMs 2, 4, 3, 5, and 7, respectively. After finding the Steiner tree rooted at VM 7, the demand is then routed to destination 8 by traversing switches 4 and 6, and directly to destination 9. The total cost in the end of the second phase is 45. ■

In the following, we present several important characteristics for graph $\mathcal{G}$, which play crucial roles to derive the approximation ratio. First, the cost of a walk $(u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$ from $s = u_1$ to the last VM $u = u_{|\mathcal{C}|+1}$ without traversing a node multiple times in $\mathcal{G}$ is equal to the sum of the total setup cost of $u_2, u_3, \cdots, u_k$, plus the total connection cost of the shortest paths between every u_j and u_{j+1} for $j = 1, 2, \cdots, k-1$ in G. Second, the edge costs in $\mathcal{G}$ satisfy the triangular inequality, as described in the following lemma. For readability, the detailed proof is presented in Appendix B.

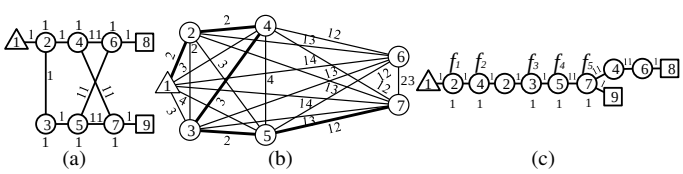

Fig. 3. Example of service overlay forest by SOFDA-SS. (a) The input network G. (b) The constructed instance of the k-Stroll problem $\mathcal{G}$, where the walk $W_{\mathcal{G}}$ between nodes 1 and 5 is shown in bold. (c) The service overlay forest with $\mathcal{C} = (f_1, f_2, f_3, f_4, f_5)$ constructed for G.

Lemma 1. *The graph $\mathcal{G}$ satisfies triangular inequality.*

Let $c(\mathcal{F}_M^{OPT})$ and $c(\mathcal{F}_E^{OPT})$ denote the setup and connection costs of the optimal service overlay forest $\mathcal{F}^{OPT}$, respectively. Based on the above two characteristics, the following theorem derives the approximation ratio of SOFDA-SS. The complete proof is presented in Appendix C.

Theorem 2. *The cost of F is bounded by $(2+\rho_{ST})c(\mathcal{F}^{OPT})$. That is, SOFDA-SS is a $(2+\rho_{ST})$-approximation algorithm for the SOF problem with one tree.*

Time Complexity Analysis. SOFDA-SS constructs $|M|$ instances of the k-stroll problem, and each of them employs the Dijkstra algorithm $|M|$ times to compute the edge costs of each instance, where $O(T_d)$ denotes the time to run the Dijkstra algorithm. Moreover, let $O(T_k)$ denote the time to solve a k-stroll instance [29], and let $O(T_s)$ represent the time to append a Steiner tree by [20]. Therefore, the overall time complexity is $O(|M|(T_d|M| + T_k + T_s))$.

V. General Case With Multiple Trees

In this section, we propose a $3\rho_{ST}$-approximation algorithm, named Service Overlay Forest Deployment Algorithm (SOFDA), for the general SOF problem with multiple sources. Different from SOFDA-SS, here we select multiple sources to exploit multiple trees for further reducing the total cost, and it is necessary to choose a different subset of destinations for each source to form a forest. In other words, both the last VMs and the set of destinations are necessary to be carefully chosen for the tree corresponding to each source. To effectively solve the above problem, our idea is to identify a short service chain from each source to each destination as a candidate service chain and then encourage more destinations to merge their service chains into a service tree, and those destinations will belong to the same tree in this case. More specifically, SOFDA first constructs an auxiliary graph $\mathbb{G}$ with each candidate service chain represented by a new virtual edge connecting the source and the last VM of the chain. Also, every source is connected to a common virtual source. SOFDA finds a Steiner tree spanning the virtual source and all destinations, and we will prove that the cost of the tree in $\mathbb{G}$ is no greater than $3\rho_{ST}c(\mathcal{F}^{OPT})$.

Nevertheless, a new challenge arises here because the service chains corresponding to the selected virtual edges in the above approach may overlap in a few nodes in G, and the solution thereby is not feasible if any overlapping node in this case needs to support multiple VNFs (see the definition of SOF). SOFDA in Section V-B thereby revises the above solution into multiple feasible trees, and we prove that SOFDA can still maintain the desired approximation in Section V-A. The pseudo code of SOFDA is presented in Appendix E (see Algorithm 2).

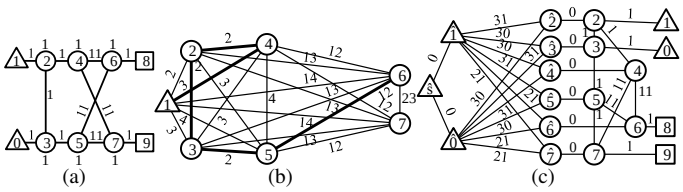

Fig. 4. Example of instance construction of the Steiner tree problem. (a) The input network G. (b) The instance of the k-Stroll problem $\mathcal{G}$ constructed, where the walk $W_{\mathcal{G}}$ between nodes 1 and 6 is shown in bold. (c) The instance of the Steiner tree problem.

A. Cost-Bounded Steiner Tree

SOFDA first constructs an auxiliary graph $\mathbb{G}$ to effectively extract multiple service chains and group the destinations. Specifically, let $\mathbb{V}_{\mathbb{S}}$ consist of the duplicate $\hat{v}$ of each source $v \in S$, and let $\mathbb{V}_{\mathbb{M}}$ contain the duplicate $\hat{v}$ of each VM $v \in M$. Therefore, $\mathbb{V} = V \cup \{\hat{s}\} \cup \mathbb{V}_{\mathbb{S}} \cup \mathbb{V}_{\mathbb{M}}$, where $\hat{s}$ denotes the *virtual source*. Also, let $\mathbb{E}_{\hat{s}\mathbb{S}}$ include the edges between $\hat{s}$ and $\hat{v}$ for each $\hat{v} \in \mathbb{V}_{\mathbb{S}}$. Let $\mathbb{E}_{\mathbb{SM}}$ consist of the *virtual edges* (representing the candidate service chain) between $\hat{v}$ and $\hat{u}$ for each $\hat{v} \in \mathbb{V}_{\mathbb{S}}$ and $\hat{u} \in \mathbb{V}_{\mathbb{M}}$, and let $\mathbb{E}_{\mathbb{M}M}$ include the edges between v and $\hat{v}$ for each $v \in M$. Then, $\mathbb{E} = E \cup \mathbb{E}_{\hat{s}\mathbb{S}} \cup \mathbb{E}_{\mathbb{SM}} \cup \mathbb{E}_{\mathbb{M}M}$. Moreover, the cost of each edge in $\mathbb{E}_{\hat{s}\mathbb{S}} \cup \mathbb{E}_{\mathbb{M}M}$ is assigned to 0, and the cost of the *virtual edge* between $\hat{v} \in \mathbb{V}_{\mathbb{S}}$ and $\hat{u} \in \mathbb{V}_{\mathbb{M}}$ in $\mathbb{E}_{\mathbb{SM}}$ is equal to the cost of the k-stroll walk that visits $|\mathcal{C}|$ VMs between v and u in G. We first present an illustrative example for the above graph transformation.

Example 3. Fig. 4 presents an example to construct the instance $\mathbb{G} = \{\mathbb{V}, \mathbb{E}\}$ of the Steiner tree problem with the graph G shown in Fig. 4(a), where $S = \{0, 1\}$, $M = \{2, 3, 4, 5, 6, 7\}$. The output $\mathbb{G}$ is presented in Fig. 4(c). SOFDA first replicates G in $\mathbb{G}$. Subsequently, it duplicates the sources 0 and 1 by creating nodes $\hat{0}$ and $\hat{1}$, and VMs $2, 3, 4, 5, 6, 7$ by creating nodes $\hat{2}, \hat{3}, \hat{4}, \hat{5}, \hat{6}, \hat{7}$ in $\mathbb{G}$. Then, the costs of edges $(\hat{s}, \hat{0})$, $(\hat{s}, \hat{1})$, $(\hat{2}, 2)$, $(\hat{3}, 3)$, $(\hat{4}, 4)$, $(\hat{5}, 5)$, $(\hat{6}, 6)$, and $(\hat{7}, 7)$ are all set to 0. To derive the cost of the virtual edge $(\hat{1}, \hat{6})$, SOFDA finds the walk from source 1 to VM 6 in G as follows. First, it constructs an instance of the k-stroll problem $\mathcal{G}$ shown in Fig. 4(b). Then, we obtain a walk $W'_{\mathcal{G}} = W_{\mathcal{G}} = (1, 4, 2, 3, 5, 6)$ in $\mathcal{G}$. By combining the shortest paths with each path connecting two consecutive nodes in $W'_{\mathcal{G}}$, we find the desired walk $W_G = (1, 2, 4, 2, 3, 5, 6)$ in G. Thus, the cost of link $(\hat{1}, \hat{6})$ is set to the cost of W_G, which is equal to $c(2) + c(4) + c(3) + c(5) + c(6) + c(1, 2) + c(2, 4) + c(4, 2) + c(2, 3) + c(3, 5) + c(5, 6) = 21$. ■

The following lemma first indicates that the cost of the constructed Steiner tree in $\mathbb{G}$ is bounded by $\rho_{ST} \cdot 3c(\mathcal{F}^{OPT})$, by showing that there is a feasible Steiner tree $\mathbb{T} = \{\mathbb{V}_{\mathbb{T}}, \mathbb{E}_{\mathbb{T}}\}$ in $\mathbb{G}$ with the cost bounded by $3c(\mathcal{F}^{OPT})$.

Lemma 2. *A feasible Steiner Tree with the cost no greater than $3c(\mathcal{F}^{OPT})$ exists in $\mathbb{G}$.*

Proof: We first show that there is a $\mathbb{T}$-like graph, $\mathbb{T}' = \{\mathbb{V}_{\mathbb{T}'}, \mathbb{E}_{\mathbb{T}'}\}$, with a cost of at most $3c(\mathcal{F}^{OPT})$ in $\mathbb{G}$. Afterward, we extract the desired $\mathbb{T}$ from $\mathbb{T}'$. Let D_v^{OPT} denote the set of the destinations in the service overlay tree rooted at source v in $\mathcal{F}^{OPT}$. In addition, for the service overlay tree rooted at source v in $\mathcal{F}^{OPT}$, let m_v^{OPT} be the representative last VM chosen from all the VMs running $f_{|\mathcal{C}|}$ on the paths from v to the destinations in D_v^{OPT}. Moreover, let T_v be the optimal Steiner tree rooted at m_v^{OPT} that spans all destinations in D_v^{OPT} in G. Then, let $\mathbb{V}_{\mathbb{T}'}$ consist of 1) $\hat{s}$, 2) the duplicate $\hat{v}$ (in $\mathbb{V}_{\mathbb{S}}$) of each source v in $\mathcal{F}^{OPT}$, 3) the duplicate $m_v^{\hat{OPT}}$ (in $\mathbb{V}_{\mathbb{M}}$) of each m_v^{OPT} in $\mathcal{F}^{OPT}$, 4) each m_v^{OPT} in $\mathcal{F}^{OPT}$, and 5) all VMs and switches (including all destinations in D) in all optimal Steiner trees T_v in G. Let $\mathbb{E}_{\mathbb{T}'}$ include the edges between 1) $\hat{s}$ and $\hat{v}$, 2) $\hat{v}$ and $m_v^{\hat{OPT}}$, 3) m_v^{OPT} and $m_v^{\hat{OPT}}$ for each spanned source v in $\mathcal{F}^{OPT}$, and 4) all links in all optimal Steiner trees T_v in G for each used source v in $\mathcal{F}^{OPT}$.

Note that for each source v in $\mathcal{F}^{OPT}$, the cost of the edge between $\hat{v}$ and $m_v^{\hat{OPT}}$ in $\mathbb{T}'$ is bounded by twice of the cost of the shortest walk that visits $|\mathcal{C}|$ VMs between v and m_v^{OPT} in G. Since there is a walk between v and v^{OPT} in $\mathcal{F}^{OPT}$, the total cost of the edges in $\mathbb{E}_{\mathbb{T}'} \cap \mathbb{E}_{\mathbb{SM}}$ is bounded by $2c(\mathcal{F}^{OPT})$. In addition, the cost of T_v is restricted by the connection cost of the service overlay tree with root v in $\mathcal{F}^{OPT}$, because the latter one not only spans m_v^{OPT} and the destinations but also spans the source v and other VMs (running $f_1, f_2, ..., f_{|\mathcal{C}|}$). Thus, the total cost of every edge in $\mathbb{E}_{\mathbb{T}'} \cap E$ is bounded by $c(\mathcal{F}^{OPT})$. Since the cost of each edge in $\mathbb{E}_{\mathbb{T}'} \cap \mathbb{E}_{\hat{s}\mathbb{S}}$ or $\mathbb{E}_{\mathbb{T}'} \cap \mathbb{E}_{\mathbb{M}M}$ is 0, the cost of $\mathbb{T}'$ is bounded by $3c(\mathcal{F}^{OPT})$. Furthermore, there is a subgraph (more specifically, a tree) $\mathbb{T}$ of $\mathbb{T}'$ that spans the virtual node and all the destinations in $\mathbb{G}$. Hence, the cost of $\mathbb{T}$ is smaller than that of $\mathbb{T}'$ and is bounded by $3c(\mathcal{F}^{OPT})$. ■

B. Cost-Bounded Service Overlay Forest

After finding a Steiner tree $\mathbb{T}$ in $\mathbb{G}$ with a bounded cost of at most $3\rho_{ST}c(\mathcal{F}^{OPT})$ by the above ρ_{ST}-approximation algorithm, to limit the total cost of the service overlay forest, SOFDA will deploy each service chain with the corresponding virtual edge in $\mathbb{T} \cap \mathbb{E}_{\mathbb{SM}}$ and the route traffic via the edges in $\mathbb{T} \cap E$. Specifically, SOFDA first 1) adds each corresponding walk of the spanned virtual edge one by one in G and then 2) adds all VMs, switches, and links in $\mathbb{T} \cap G$ to F.

Example 4. Fig. 6 presents an example for the construction of the service overlay forest with $\mathcal{C} = (f_1, f_2, f_3, f_4, f_5)$ in Fig. 4(a) by SOFDA. First, an instance $\mathbb{G} = \{\mathbb{V}, \mathbb{E}\}$ of the Steiner Tree problem is constructed with the input parameters G, $S = \{0, 1\}$, $M = \{2, 3, 4, 5, 6, 7\}$, and $\mathcal{C} = (f_1, f_2, f_3, f_4, f_5)$, and a Steiner tree $\mathbb{T}$ in $\mathbb{G}$ using the ρ_{ST}-approximation algorithm in [20] is obtained, as shown in Fig. 6(a). ■

Nevertheless, multiple walks in G corresponding to the spanned virtual edges in $\mathbb{T}$ may overlap in a few VMs, and the solution in this case is infeasible if any overlapping VM in this case needs to perform different VNFs (see the definition of SOF in Section III). The situation is called *VNF conflict* in this paper. In the following, we present an effective way to eliminate the conflict by tailoring the overlapping walks without increasing the cost. To address the VNF conflict, when a walk $W_G = (v_1, v_2, \cdots, v_n)$ in G is added to the service overlay forest F, it is encouraged to augment F with a modified walk $W = (u_1, u_2, \cdots, u_n)$ based on W_G. Note that a VM or switch is allowed to be passed without processing any VNF by simply forwarding the data. Moreover, a VNF conflict happens when two walks compete for a clone to perform different VNFs. Fig. 6(b) presents an example of the VNF conflict, where W_1 and W_2 respectively desire to run f_1 and f_4 on VM 4. Suppose that a walk W (between source s and VM v) faces the VNF conflict with another walk W_1 (between source

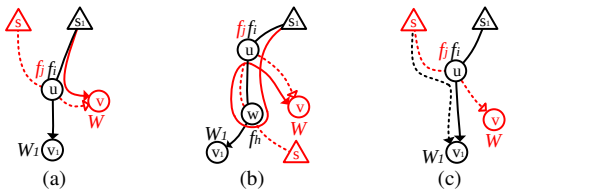

Fig. 5. Resolve of *VNF conflicts* between two walks, where the black solid (or dashed) line denotes the original (or updated) W_1, and the red dashed (or solid) line denotes the original (or updated) W_2. Fig. 5 (a), (b), and (c) show the resolve of the first, second, and third kinds of *VNF conflicts*, respectively.

s_1 and VM v_1) in F. We solve the conflict between W and W_1 by changing the source of W from s to s_1 (attaching W to W_1), or changing the source of W_1 from s_1 to s (attaching W_1 to W) without adding new links, VMs, and switches to F and without enabling new VMs in F for VNFs.

Example 5. Following Example 4, SOFDA finds walks $W_{G,1} = (1, 2, 4, 2, 3, 5, 6)$ and $W_{G,2} = (0, 3, 5, 3, 2, 4, 7)$ in G, where f_1, f_2, f_3, f_4, f_5 are installed at VMs 4, 2, 3, 5, 6 on $W_{G,1}$, and also at VMs 3, 5, 2, 4, 7 on $W_{G,2}$, respectively. After $W_{G,1}$ is added to F, we have $F = \{W_1\}$, where W_1 consists of one clone for source 1, two clones of VM 2, and one clone for VMs 4, 3, 5, 6 due to $F = \emptyset$ in the beginning. As $W_{G,2}$ is added to F, SOFDA augments F with W_2, where W_2 includes one clone for source 0, one clone for VMs 3, 5, 2, 4 (on which f_3, f_4, f_2, and f_1 are already running on W_1), and two new clones for VMs 3 and 7, as shown in Fig. 6(b). ■

Specifically, let u be the first VM, where W experiences the VNF conflict with W_1 by backtracking W. Recall in Fig. 6(b), for example, that VM 4 is the first conflict node with W_1 by backtracking W. Let $f_1, f_2, \cdots, f_{|\mathcal{C}|}$ denote the VNFs required to be performed in sequence on W and W_1. Let f_i and f_j be the VNFs located at u on W_1 and W, respectively. SOFDA effectively addresses the VNF conflict in details as follows.

First, if $j \leq i$, SOFDA attaches W to W_1 through u by changing W to the concatenation of the sub-walk of W_1 from s_1 to u (on which $f_1, f_2, \cdots, f_i$ are installed in sequence, identical to W_1) and the sub-walk of W from u to v (on which $f_{i+1}, f_{i+2}, \cdots, f_{|\mathcal{C}|}$ are running in sequence, identical to W), as shown in Fig. 5(a).

Example 6. Following Example 5, W_2 first experiences the *VNF conflict* with W_1 at (the clone of) VM 4, where f_4 and f_1 are installed on W_2 and W_1, respectively. The sequence numbers of the VNFs at VM 4 on W_2 and W_1 are 4 and 1, respectively. The condition $j \leq i$ is not satisfied since $j = 4$ and $i = 1$. SOFDA then checks the next condition. Note that one of the three conditions must be satisfied. ■

Second, if there is another VM w such that W experiences the *VNF conflict* with W_1 at w, where f_h with $h \geq j$ is on W_1, SOFDA attaches W to W_1 through w by changing W to the concatenation of the sub-walk of W_1 from s_1 to w (on which $f_1, f_2, \cdots, f_h$ are running in sequence, identical to W_1), the sub-walk of W from w to u, and the sub-walk of W from u to v (on which $f_{h+1}, f_{h+2}, \cdots, f_{|\mathcal{C}|}$ are running in sequence, identical to W), as illustrated in Fig. 5(b).

Example 7. Following Example 6, W_2 experiences another *VNF conflict* with W_1 at VM 5, where f_2 and f_4 are performed on W_2 and W_1, respectively. Since the sequence number of the

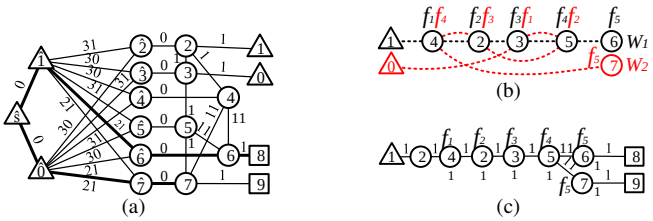

Fig. 6. Example of construction of the service overlay forest by SOFDA. (a) The Steiner tree, shown in bold line. (b) Two walks with *VNF conflict*. (c) The service overlay forest with $\mathcal{C} = (f_1, f_2, f_3, f_4, f_5)$ constructed for G in Fig. 4(a).

VNF at VM 5 on W_1 is not smaller than that of the VNF at VM 4 on W_2, SOFDA attaches W_2 to W_1 through VM 4 as follows. SOFDA first steers W_2 along the sub-walk of W_1 from source 1 to VM 5 (i.e., the walk $(1, 2, 4, 2, 3, 5)$) on which f_1, f_2, f_3, f_4 are running in sequence at VMs 4, 2, 3, 5, respectively, identical to W_1. Subsequently, it continues steering W_2 along the sub-walk of W_2 from VM 5 to VM 4 (i.e., the walk $(5, 3, 2, 4)$), and the sub-walk of W_2 from VM 4 to VM 7 (i.e., the walk $(4, 7)$) on which f_5 is run at VM 7, identical to W_2. Finally, the sub-walk $(5, 3, 2, 4, 7)$ on the revised W_2 can be shortened to be a walk $(5, 7)$. The constructed service overlay forest for G is displayed in Fig. 6(c). ■

Otherwise, SOFDA attaches W_1 to W through u by changing W_1 to the concatenation of the sub-walk of W from s to u (on which $f_1, f_2, \cdots, f_j$ are running in sequence, identical to W) and the sub-walk of W_1 from u to v_1 (on which $f_{j+1}, f_{j+2}, \cdots, f_{|\mathcal{C}|}$ are run in sequence, identical to W_1), as shown in Fig. 5(c). Moreover, when a walk W experiences the *VNF conflict* with multiple walks $W_1, W_2, \cdots, W_l$ in F in sequence by backtracking W, SOFDA resolves the *VNF conflict* between W and $W_1, W_2, \cdots, W_l$ one-by-one. The following theorem derives the approximation ratio for SOFDA.

Theorem 3. *The cost of the constructed service overlay forest F is bounded by $3\rho_{ST}c(\mathcal{F}^{OPT})$.*

Proof: First, the cost of Steiner tree $\mathbb{T}$ in $\mathbb{G}$ is bounded by ρ_{ST} times of the optimal Steiner tree in $\mathbb{G}$. Since the cost of the optimal Steiner tree in $\mathbb{G}$ is bounded by $3c(\mathcal{F}^{OPT})$ according to Lemma 2, the cost of $\mathbb{T}$ is limited by $3\rho_{ST}c(\mathcal{F}^{OPT})$. In addition, since the cost of the edge between $\hat{v} \in \mathbb{V}_{\mathbb{S}}$ and $\hat{u} \in \mathbb{V}_{\mathbb{M}}$ of u in $\mathbb{T}$ is identical to the cost of the walk that visits $|\mathcal{C}|$ VMs between v and u in G, the cost of F constructed in G is equal to the cost of $\mathbb{T}$ and thereby bounded by $3\rho_{ST}c(\mathcal{F}^{OPT})$ if no *VNF conflict* occurs in F. On the other hand, when the *VNF conflict* between two walks happens, one of the two walks in F is updated, and no new link, VM, and switch is added to F, and no VM in F is newly created to perform the VNF. Thus, the cost of F revised for resolving the *VNF conflict* is still bounded by $3\rho_{ST}c(\mathcal{F}^{OPT})$. The theorem follows. ■

Time Complexity Analysis. We follow the notations in the time complexity analysis of SOFDA-SS. To generate the instance of the Steiner tree problem, SOFDA constructs $|S||M|$ instances of the k-stroll problem, and each of them employs the Dijkstra algorithm $|M|$ times to compute the edge costs of each instance. Then, SOFDA solves the k-stroll instance by [29] to derive the costs of virtual edges (i.e., corresponding candidate service chains). To eliminate the conflict, in the worst case, all the added walks in F are appended to the newly added walk, and the complexity is $O(|M|^3)$. There-

fore, the total time complexity is dominated by constructing and solving k-stroll instance and finding a Steiner tree, i.e., $O(|S||M|(|M|T_d + T_k) + T_s)$.

VI. Distributed Implementation

For large SDNs, it is important to employ multiple SDN controllers, where each one monitors and controls a subset of the network [30]–[32], and the communication protocols [33] between controllers are developed to facilitate scalable control of the whole network. In the following, therefore, we discuss the distributed implementation of the proposed algorithm in Section V to support multi-controller SDNs. Note the controller that receives the request is elected to be the *leader*, which is responsible for progress tracking and phase switching.

First, shortest-path routing plays a fundamental role in SOFDA to build the auxiliary graph $\mathbb{G}$ and the service chain corresponding to each edge in $\mathbb{G}$. To find a shortest path traversing multiple domains, it is necessary for each controller to first abstract a matrix that consists of the lengths between every pair of border routers over the Southbound interface [33] within its domain. Afterward, each controller propagates the matrix to the other controllers along with the Network Layer Reachability Information of SDNi Wrapper over East-West Interface. which is used to share the connectivity information with the neighboring controllers. More specifically, let s and t denote the source and the destination, respectively. The controller C_s covering s can find the corresponding domain by the IP prefix of t. Then, controller C_s informs the controller C_t that covers t of the lengths of all shortest paths from s to all broader routers of C_t. Afterward, controller C_t can respond the best broader router to controller C_s, and the length of a shortest path can be acquired accordingly.

Equipped with the shortest-path computation from multiple controllers, each controller can acquire the length of each shortest path between a VM in its domain and any other VM (or source). Thus, once the forest construction is initiated, every controller that covers a source will communicate with other controllers to collect the matrices of lengths between any two VMs and the lengths between any source and any VM. Then, the controller can find all candidate service chains from its covered source to each VM and creates a virtual link in $\mathbb{G}$ representing the service chain to connect the virtual source and the corresponding last VM.

Afterward, a distributed Steiner tree algorithm [34] can be employed by multiple controllers to find the Steiner tree, where the computation load originally assigned to each switch in the distributed algorithm can be finished by its controller instead. In SOFDA, it is important to address the VNF conflicts in multiple domains. To achieve this goal, each controller first removes the *useless* candidate service chains that do not connect with any destination, and then informs any other controller whose coverage is visited by any remaining service chain. When one of the informed controllers observes a VNF conflict of two service chains, it notifies the other controller to collaboratively remove the conflict according to the conflict elimination algorithm described in Section V-B. Finally, each controller deletes the virtual source, deploys the remaining service chains, and forwards the content to the destinations by SOF.

VII. Discussion

A. Static Mulitcast Trees with Service Chaining

To the best knowledge of the authors, this paper is the first one that explores the notion of the *service forest*, i.e., the fundamental multi-tree multicast problem with service chaining, and provides approximation algorithms with theoretical bounds. Therefore, we first consider the fundamental problems for static SDN/NFV multicast and then extend the proposed algorithms to the dynamic case in Section VII-C.

Actually, static multicast is crucial for backbone ISPs. In this situation, each terminal node of a multicast tree is usually an edge router or a local proxy server of the ISP, instead of a dynamic user client. For example, current live streams are sent by the source (i.e., headends or content servers) and travel through the high-speed backbone network to the access nodes and edge nodes (e.g., Digital Subscriber Line Access Multiplexer (DSLAM) [35], or a Mobile Edge Computing (MEC) server [36]) via *static multicast* trees [17], [35]–[39] (e.g., Chung-Hwa Telecom MOD [40]), whereas the dynamic user join and leave are handled by the local access nodes and edge nodes. Static multicast trees can significantly reduce the backbone bandwidth consumption for each stream and thereby is much more scalable to support a large number of video channels. In this case, each access node usually serves hundreds or thousands of end users and streams one (or few) channel(s) to each user according to the available bandwidth between the access node and user devices (e.g., set-top boxes). Moreover, for the massively multi-user virtual reality (e.g., gaming) [41]–[45], the servers create a virtual environment with a 3D model, player avatars, and scripts, and then transmit the data by static multicast to several MEC servers [36], which always need to appear in a multicast group. In the above life examples, our proposed algorithms can facilitate static multicast with service chaining (i.e., multiple stages of servers) to support a large number of streams between the headend server and the local access nodes/edge nodes.

B. Cost Model and Online Deployment

In the online scenario, when a new request arrives, SOFDA allocates the required resources for the request by constructing a service forest according to the current link and node costs. To balance the network resource consumption and accommodate more requests in the future, congested links and nodes are unnecessary to be assigned with higher costs for encouraging SOFDA to employ the links and nodes with low loads [46]–[48]. In this paper, therefore, we exploit [46], which is designed for online adaptive routing in the Internet, to assign a convex cost to each link or node. The cost will significantly increase as the load linearly grows, to avoid overwhelming the link or node. More specifically, let l and p denote current load and capacity of the link or node, respectively, and the cost c is set according to the utilization (i.e., l/p) as follows and illustrated in Fig. 7.

$$c = \begin{cases} l & \text{if } l/p \leq 1/3, \\ 3l - 2/3p & \text{else if } l/p \leq 2/3, \\ 10l - 16/3p & \text{else if } l/p \leq 9/10, \\ 70l - 178/3p & \text{else if } l/p \leq 1, \\ 500l - 1468/3p & \text{else if } l/p \leq 11/10, \\ 5000l - 14318/3p & \text{otherwise.} \end{cases}$$

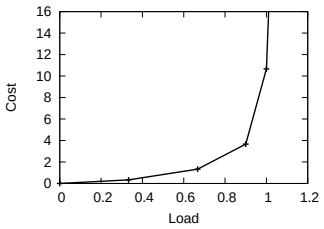

Fig. 7. The cost function with different load l and capacity $p = 1$.

The cost model properly handles the online situation by assigning a huge cost to a more congested node or link. Therefore, SOFDA will avoid choosing the above congested node or link to minimize the total cost of the service forest. SOFDA thereby can mitigate the impact on a VM of other VMs colocated with an overloaded node. Indeed, the cost model can be applied to both private and public cloud networks, where resource optimization and load balancing are usually addressed. For example, Chung-Hwa Telecom MOD [40] is built in its private clouds while Netflix [8] adopts AWS [9].

Nevertheless, each request has a different duration, and an approach without considering the duration of the request is inclined to incur fragmentation of the network resources and degrade the performance. However, the duration of a stream (e.g., a VR multi-player game) is usually difficult to be precisely predicted, and many current approaches thereby adaptively reroute [49]–[53] and migrate the VM [48], [54]–[57] to relocate the network resources when congestion occurs. Similar to the above approaches, when a node or link becomes congested, SOFDA reroutes the service forest by letting the users downstream to the above node or link re-join the forest again (explained in the reply of the first question), where the current path in the forest is removed only after the new join path is created to avoid service interruption [49]–[53].

C. Adjustments for Various Dynamic Cases

. In the following, we extend SOFDA to support the dynamic join and leave of destination users and the addition and deletion of NFVs in a service forest after a session starts. To address the dynamic case, a simple approach is to run SOFDA again for the whole forest. Nevertheless, this approach tends to incur massive computation loads in the SDN controller, especially when users frequently join and leave the multicast group or change the computation tasks in the service forest. In the following, therefore, we extend SOFDA to properly handle the dynamic case [21], [22].

1) **Destination Leave.** When a destination v leaves the service forest, if v is a leaf node, SOFDA removes v and all intermediate nodes and links in the path connecting v and the closest upstream branch node in the service forest, where a branch node is a node in the forest with at least two child nodes. By contrast, if v is not a leaf node, because there are other destination users in the subtree rooted at v, SOFDA is not allowed to remove the path connecting to the upstream branch node.
2) **Destination Join.** When a new destination user v joins the service forest, SOFDA finds the walk from v to the forest with the lowest cost. More specifically, for each node u in the forest $\mathcal{F}$ that can be a candidate branch node to connect v, let $f(u)$ denote the index of the last installed VNF between a source s and u in the forest. To derive the cost in the walk from u to v, SOFDA finds the walk with $k = |\mathcal{C}| - f(u) + 1$ from u to v to install the $(|\mathcal{C}| - f(u))$ new VNFs in the walk by exploiting k-stroll in the transformed graph (see Section IV). Let $W_G(u, v) = (u_1, ..., u_k)$ denote the acquired walk, where $u_1 = u$ and $u_k = v$. In this case, SOFDA needs to install the new VNFs $f_{f(u)+1}, ..., f_{|\mathcal{C}|}$ on the above walk from u to v, and the cost of the forest is increased by $\min_{u \in \mathcal{F}}\{c(W_G(u, v))\}$. SOFDA carefully examines every possible u in the existing service forest to effectively minimize the increasing cost, and the node u leading to the smallest cost is selected to serve the new destination user v accordingly.
3) **VNF Deletion.** When VNF f_j is removed from the service forest, for each VM v that installs a VNF f_j, SOFDA connects the VM u with the upstream VNF f_{j-1} to the VM w with the downstream VNF f_{j+1} (along the minimum-cost path from u to w in the original G) in the forest, where the source (or destination) can be regarded as the VM with the upstream (or downstream) VNF f_{j-1} (or f_{j+1}) if f_j is the first (or last) VNF.
4) **VNF Insertion.** When VNF f_j is inserted to the service forest, for each pair of VMs u and w with VNFs f_{j-1} and f_{j+1}, respectively, SOFDA installs f_j on an available VM v, and connects u to v and v to w in the forest such that the sum of 1) the connection cost of the path between u and v, 2) the setup cost of v, and 3) the connection cost of the path between v and w is minimized. When f_j is the first (or last) VNF, the source (or destination) is regarded as the VM with VNF f_{j-1} (or f_{j+1}). In addition, if two pairs of VMs (u_1, w_1) and (u_2, w_2) with VNFs f_{j-1} and f_{j+1} choose the same VM v to install VNF f_j, SOFDA removes all intermediate nodes and links in the path connecting u_2 and v in the forest in order to reduce the total cost of the forest (i.e., avoid creating redundant paths in the forest).
5) **Link Congestion.** For any congested link e between the VMs with VNF f_j and f_{j+1}, SOFDA updates the link cost according to [46] and then re-connects the two VMs with the path associated with the lowest cost. SOFDA can effectively avoid choosing a congested link because the cost of the link will be extremely large. On the other hand, if e is between the source and a VM (or VM and a destination), the source (or the destinations) is regarded as the upstream VM (or the downstream VM) and handled in a similar manner.
6) **VM Overload.** For any overloaded VM v between the VMs with VNF f_{j-1} and f_{j+1}, SOFDA updates the node cost according to [46] to find an available VM v' and then re-connects it to the upstream VM and downstream VM with the path having the lowest cost. Therefore, SOFDA can also avoid selecting an overloaded VM. On the other hand, if v is the first VNF (or the last VNF), the source (or the destinations) is regarded as the upstream VM (or the downstream VM) and then handled in a similar manner.

VIII. Numeric Result

A. Simulation Setup

We conduct simulations to compare different approaches in two inter-data-center networks: IBM SoftLayer [58] and Cogent [59]. SoftLayer contains 27 access nodes with 49 links and 17 data centers, whereas Cogent has 190 access nodes with

260 links and 40 data centers. We also generate a synthetic network with 5000 access nodes, 10000 links , and 2000 data centers by Inet [60]. The edge costs and the node costs are set according to [46] and [48] based on the corresponding loads (described in Section VII-B), respectively. The sources and destinations are chosen uniformly at random from the nodes in the network. We examine the performance of different approaches in two scenarios: one-time deployment and online deployment. Moreover, we also implement all algorithms in a small-scale SDN with HP SDN switches.

In the one-time deployment scenario, the link bandwidth is set to 100 Mbps, and each requested demand is set to 5 Mbps. The link usage is randomly chosen in $(0, 1)$ so as to derive the edge cost according to [46]. Also, the total number of VMs ranges in $\{5, 15, 25, 35, 45\}$, and each VM is randomly attached to a data center. The service chain length (i.e., the number of VNFs in the chain) ranges in $\{3, 4, 5, 6, 7\}$ in both Cogent and Softlayer networks. The number of destinations and candidate sources range in $\{2, 4, 6, 8, 10\}$ and $\{2, 8, 14, 20, 26\}$, respectively in both networks. The default numbers of candidate sources, destinations, VMs, and service chain length are 14, 6, 25, 3, respectively.

Afterward, for the online deployment scenario, the node/link usages are zero initially. Each data center has 5 VMs with the cost according to the host machine utilization. Afterward, we incrementally generate a new request, and the node costs and edge costs will be updated according to [46]. The numbers of destinations and candidate sources in the request are randomly chosen from 13 to 17 and 8 to 12 in Softlayer, and they are from 20 to 60 and from 10 to 30 in Cogent, respectively. The number of demanded services in a request is 3.

We compare the proposed algorithm with the following ones. 1) CPLEX [61]. It finds the optimal solution by solving the IP formulation in Section III-A. 2) Enhanced Steiner Tree (eST). Since the Steiner tree algorithm [20] does not select VMs in the tree, we extend it for SOF as follows. We find the minimum-cost tree among all Steiner trees rooted at different sources. Afterward, we construct the shortest service chain that is closest to the tree from [13], [62] and then connect it to the tree with the minimum cost. 3) Enhanced algorithm for the NFV enabled multicast problem (eNEMP). Since the algorithm for the NFV enabled multicast problem (NEMP) [27] does not support multiple sources and VNFs, similar to the above extension, we construct a service chain and then connect it to the tree, where the chain spans the VM that has been chosen in the tree. Moreover, we enable eST and eNEMP to support multiple sources via the modification as follows. The idea is to iteratively add a service tree in the solution until no tree can reduce the total cost. At each iteration, we elect the minimal-cost service tree among all candidate trees rooted at each unused source, run VNFs sequentially on unused VMs, and span all the destinations in D. To estimate the profit of tree addition, we calculate the total cost of the current forest with the elected tree, where each destination is spanned and served by the closet tree. Hence, we add the elected tree and proceed to the next iteration if it can decrease the total cost. Otherwise, we output the forest. Furthermore, a special case with only one Steiner tree connected with a service chain (denoted by ST in the figures) is also evaluated.

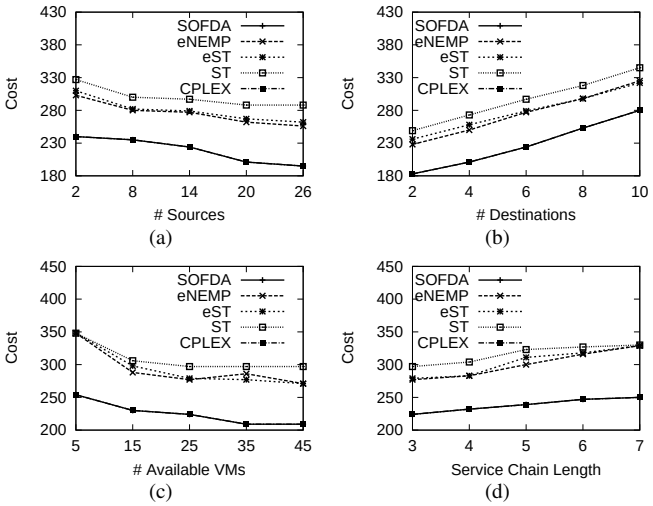

Fig. 8. The impact on cost of (a) the number of sources, (b) the number of destinations, (c) the number of available VMs, and (d) the service chain length in Softlayer network.

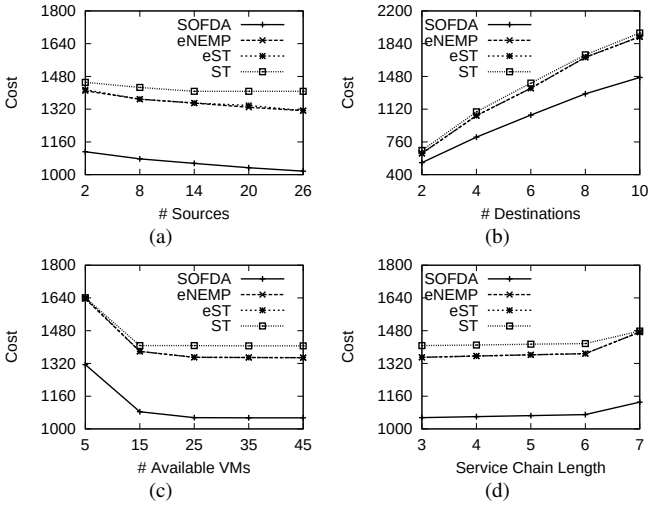

Fig. 9. The impact on cost of (a) the number of sources, (b) the number of destinations, (c) the number of available VMs, and (d) the service chain length in Cogent network.

B. One-Time Deployment

We compare the performance of SOFDA, eNEMP, eST, ST, and the optimal solution generated by CPLEX with different numbers of sources, destinations, VMs, and different numbers of demand services. Because SOF is NP-hard, CPLEX is able to find the optimal solutions for small instances, and thus only Softlayer is tested in this case. Figs. 8 and 9 manifest that SOFDA is very close to the optimal solutions, and choosing multiple sources effectively reduces the total cost. The improvement in Fig. 9 and Fig. 10 is more significant because larger networks (i.e., Cogent and the synthetic network) contains more candidate nodes and links to generate a more proper forest. Since eNEMP and eST do not choose multiple sources and VMs during the multicast routing, they tend to miss many good opportunities for allocating the VMs with small costs to the tree with fewer edges. By contrast, the results indicate that SOFDA effectively reduces the total cost by 30%. Also, when the number of sources increases,

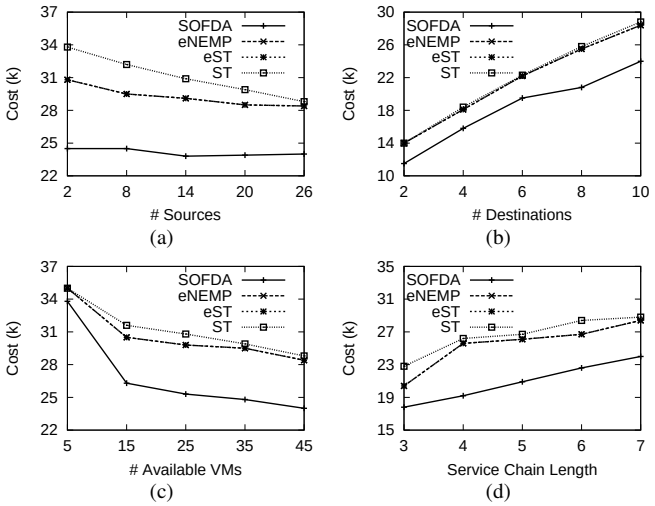

Fig. 10. The impact on cost of (a) the number of sources, (b) the number of destinations, (c) the number of available VMs, and (d) the service chain length in the synthetic network by Inet.

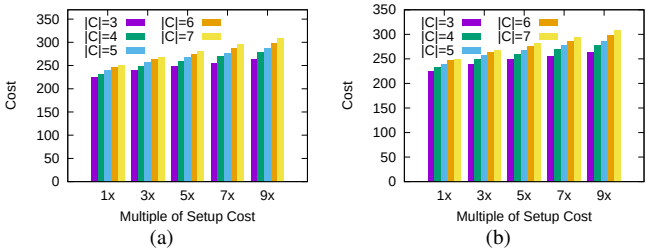

Fig. 11. The impact on (a) cost and (b) average number of used VMs by different multiples of setup cost and service chain length.

the destinations have more candidate trees to join, and thus the total cost is effectively reduced. However, the total cost increases when the number of destinations grows, because a service tree is necessary to span more destinations. Fortunately, when we have more VMs, there are more candidate machines to deploy VNFs, and the total cost thereby can be reduced. Fig. 11 presents the impact of different setup costs. The forest cost increases as the setup cost (i.e., 1x, 3x, ..., 9x) or the length of a demanded service chain (i.e., $|\mathcal{C}|$) grows as shown in Fig. 11(a). Fig. 11(b) manifests that the average number of selected VMs in a forest is effectively reduced by SOFDA as the setup cost of a VM increases. Moreover, when the length of a demanded service chain (i.e., $|\mathcal{C}|$) becomes larger, the number of required VM needs to increase in order to satisfy new user requirements.

Table I shows the running time of SOFDA with different numbers of sources and network sizes. The running time is less than 2 seconds for small networks, such as the one with 1000 nodes and 2 sources. With $|S|$ and $|V|$ increase, the running time grows, but SOFDA only requires around 19 seconds for the largest case.

C. Online Deployment

In the following, we explore the online scenario with the requests arriving sequentially. The edge costs also grow incrementally due to more traffic demand. Fig. 12 presents the accumulative costs (i.e., the total cost from the beginning to the current time slot) of different approaches. It manifests that

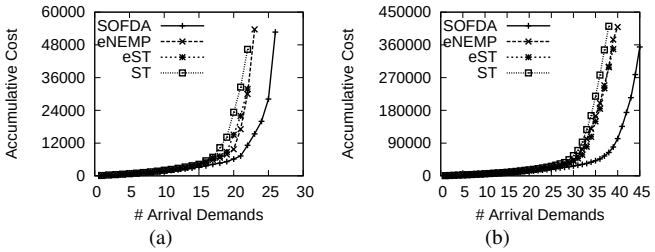

Fig. 12. Performance of the online deployment in (a) Softlayer network, and (b) Congent network.

TABLE I
THE RUNNING TIME OF SOFDA (SECONDS)

$\|V\|$	$\|S\|=2$	$\|S\|=8$	$\|S\|=14$	$\|S\|=20$	$\|S\|=26$
1000	1.35	5.15	9.08	13.16	16.03
2000	1.48	5.622	9.76	13.77	17.172
3000	1.76	5.84	9.84	14.26	18.81
4000	1.89	6.15	9.88	14.86	18.9
5000	2.25	6.87	10.99	15.75	19.65

SOFDA outperforms the others because the existing approaches focus on minimizing the traditional tree cost and thus tend to miss many good opportunities to deploy the VNFs on a longer path with sufficient VMs. By contrast, SOFDA carefully examines the edge costs and node costs and acquires the best trade-off between utilizing more VMs (leading to a smaller forest) and reducing the number of VMs, especially when the network load increases.

D. Implementation

To evaluate SOF in real environments, we implement SOFDA in Emulab [63]. The version and build of the Emulab are 4.570 and 03/17/2017, respectively. We create the topology by using NS format defined by Emulab and run Ubuntu 14.04 in each end host. We also deploy an experimental SDN with HP Procurve 5406zl OpenFlow-enabled switches and HP DL380G8 servers, where OpenDaylight is the OpenFlow controller, and OpenStack is employed to manage VMs. To support distributed computation, we run multiple OpenDaylight instances in VMs deployed in different servers and leverage the ODL-SDNi architecture [33], which enables inter-controller communications. In addition, SOFDA is implemented as an application on the top of OpenDaylight and relies on OpenDaylight APIs to install forwarding rules into the switches. It also calls OpenStack APIs to launch VM instances, which are enabled VNFs. The goal is to evaluate the transcoded and watermarked video performance under the environment with limited resources. Our testbed includes 14 nodes and 20 links, where the link capacity is set as 50 Mbps, and each node can support one VNF as explained in Section III. Two nodes are randomly selected as the video sources connecting to Youtube, and the full-HD test video is in 137 seconds encoded by H.264 with the average bit rate as 8 Mbps. Four nodes are randomly selected as destinations playing the videos with the VLC player. The video streams are processed by VNFs, including a video transcoder and a watermarker implemented by FFmpeg before reaching the destinations. The available bandwidth of each link ranges from 4.5 Mbps to 9 Mbps to emulate the scenario with the network congestion, where the video playback may stall and wait for startup again or re-buffering. During the video playback, we measure the startup

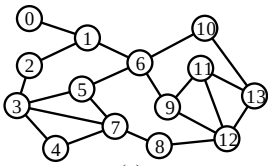

(a)

Fig. 13. Topology of Experimental SDN.

TABLE II
EVALUATIONS IN OUR EXPERIMENTAL SDN NETWORK AND EMULAB

Algorithms	Startup Latency		Re-buffering Time	
	Ours	Emulab	Ours	Emulab
SOFDA	7.5 s	5.5 s	34.0 s	29.8 s
eNEMP	9.0 s	5.9 s	39.5 s	39.0 s
eST	10.0 s	6.2 s	41.0 s	45.7 s

latency and total video re-buffering time, which are crucial for user QoE. Table II summarizes the average startup latency and re-buffering time of different approaches. The results manifest that the startup latency of SOFDA is 20 % and 33 % shorter than eNEMP and eST, and the video stalling time of SOFDA is 16% and 21% smaller than eNEMP and eST, respectively. The experiment results also indicate that SOFDA routes traffic to less congested links compared with eNEMP and eST, and fewer packets thereby are lost.

IX. Conclusion

In this paper, we investigated a new optimization problem (i.e., SOF) for cloud SDN. Compared with previous studies, the problem is more challenging because both the routing of a forest with multiple trees and the allocation of multiple VNFs in each tree are required to be considered. We proposed a $3\rho_{ST}$-approximation algorithm (SOFDA) to effectively handle the VNF conflict, which has not been explored by previous Steiner Tree algorithms. We also discussed the distributed implementation of SOFDA. Simulation results manifest that SOFDA outperforms the existing ones by over 25%. Implementation results indicate that SOF can significantly improve the QoE of the Youtube traffic. Since current IP multicast supports dynamic group membership (i.e., each user can join and leave a tree at any time), our future work is to explore the online problem for rerouting of the forest and relocation of VNFs in cloud SDN with a performance guarantee.

References

[1] "Setting the future media services architecture," *Ericsson Review*, 2015.
[2] Anvato. [Online]. Available: http://www.anvato.com/
[3] Y. Li and W. Cai, "Consistency-aware zone mapping and client assignment in multi-server distributed virtual environments," *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, pp. 1570–1579, 2015.
[4] H. Hu, R. W. H. Lau, H. Hu, and B. Wah, "On view consistency in multi-server distributed virtual environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, pp. 1428–1440, 2014.
[5] J. Marsh, M. Glencross, S. Pettifer, and R. Hubbold, "A network architecture supporting consistent rich behavior in collaborative interactive applications," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, pp. 405–416, 2006.
[6] "The need for a new approach to broadcast production," Cisco White Paper.
[7] "Network functions virtualisation - introductory white paper," ETSI White Paper.
[8] Netflix. [Online]. Available: https://www.netflix.com/
[9] AWS. [Online]. Available: https://aws.amazon.com/
[10] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. D. Turck, and R. Boutaba, "Network function virtualization: State-of-the-art and research challenges," *IEEE Communications Surveys Tutorials*, vol. 18, pp. 236–262, 2016.
[11] T. Lukovszki and S. Schmid, "Online admission control and embedding of service chains," in *SIROCCO*, 2015.
[12] M. Xia, M. Shirazipour, Y. Zhang, H. Green, and A. Takacs, "Network function placement for NFV chaining in packet/optical datacenters," *Lightwave Technology, Journal of*, vol. 33, pp. 1565–1570, 2015.
[13] T.-W. Kuo, B.-H. Liou, K. C.-J. Lin, and M.-J. Tsai, "Deploying chains of virtual network functions: On the relation between link and server usage," in *IEEE INFOCOM*, 2016.
[14] "TV and media 2016," Ericsson Consumer and Industry Insight Report.
[15] Crackle. [Online]. Available: http://www.crackle.com/
[16] Pluto tv. [Online]. Available: http://corporate.pluto.tv/
[17] R. Malli, X. Zhang, and C. Qiao, "Benefits of multicasting in all-optical networks," in *All-Optical Networking: Architecture, Control, and Management Issues*, 1998.
[18] V. Aggarwal, R. Caldebank, V. Gopalakrishnan, R. Jana, K. K. Ramakrishnan, and F. Yu, "The effectiveness of intelligent scheduling for multicast video-on-demand," in *ACM MM*, 2009.
[19] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol independent multicast-sparse mode (PIM-SM): Protocol specification," 1998, protocol specification.
[20] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità, "An improved LP-based approximation for Steiner Tree," in *ACM STOC*, 2010.
[21] S.-H. Shen, L.-H. Huang, D.-N. Yang, and W.-T. Chen, "Reliable multicast routing for software-defined networks," in *IEEE INFOCOM*, 2015.
[22] L.-H. Huang, H.-C. Hsu, S.-H. Shen, D.-N. Yang, and W.-T. Chen, "Multicast traffic engineering for software-defined networks," in *IEEE INFOCOM*, 2016.
[23] V. Aggarwal, R. Caldebank, V. Gopalakrishnan, R. Jana, K. K. Ramakrishnan, and F. Yu, "The effectiveness of intelligent scheduling for multicast video-on-demand," in *ACM MM*, 2009.
[24] H. Ma, G. K. Shin, and W. Wu, "Best-effort patching for multicast true VoD service," *Multimedia Tools and Applications*, vol. 26, pp. 101–122, 2005.
[25] M. Hofmann and L. Beaumont, *Content Networking: Architecture, Protocols, and Practice*. Morgan Kaufmann, 2005.
[26] L. H. Huang, H. J. Hung, C. C. Lin, and D. N. Yang, "Scalable and bandwidth-efficient multicast for software-defined networks," in *IEEE GLOBECOM*, 2014.
[27] S. Q. Zhang, Q. Zhang, H. Bannazadeh, and A. Leon-Garcia, "Network function virtualization enabled multicast routing on SDN," in *IEEE ICC*, 2015.
[28] M. Zeng, W. Fang, J. J. P. C. Rodrigues, and Z. Zhu, "Orchestrating multicast-oriented NFV trees in inter-DC elastic optical networks," in *IEEE ICC*, 2016.
[29] K. Chaudhuri, B. Godfrey, S. Rao, and K. Talwar, "Paths, trees, and minimum latency tours," in *IEEE FOCS*, 2003.
[30] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "SDX: A software defined internet exchange," in *ACM SIGCOMM*, 2014.
[31] S. Schmid and J. Suomela, "Exploiting locality in distributed SDN control," in *HotSDN*, 2013.
[32] G. Hampel, M. Steiner, and T. Bu, "Applying software-defined networking to the telecom domain," in *INFOCOM WKSHPS*, 2013.
[33] Inter SDN controller communication (SDNi). OpenDaylight Summit 2015. [Online]. Available: http://events.linuxfoundation.org/sites/events/files/slides/ODL-SDNi_0.pdf
[34] M. Bezenšek and B. Robič, "A survey of parallel and distributed algorithms for the steiner tree problem," *Int. J. Parallel Program.*, vol. 42, pp. 287–319, 2014.
[35] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching television over an IP network," in *ACM SIGCOMM conference on Internet measurement*, 2008.
[36] "Mobile-edge computing (mec); service scenarios," eTSI GS MEC-IEG 004. [Online]. Available: http://www.etsi.org/deliver/etsi_gs/MEC-IEG/001_099/004/01.01.01_60/gs_MEC-IEG004v010101p.pdf
[37] F. M. Ramos, R. J. Gibbens, F. Song, P. Rodriguez, J. Crowcroft, and I. H. White, "Reducing energy consumption in IPTV networks by selective pre-joining of channels," in *ACM SIGCOMM Workshop on Green Networking*, 2010.
[38] D. A. G. Manzato and N. L. S. da Fonseca, "A survey of channel switching schemes for IPTV," *IEEE Communications Magazine*, vol. 51, pp. 120–127, 2013.
[39] L. C. Yeh, C. S. Wang, C. Y. Lin, and J. S. Chen, "An innovative application over communications-asa-service: Network-based multicast IPTV audience measurement," in *APNOMS*, 2011.
[40] CHT-MOD. [Online]. Available: http://mod.cht.com.tw/
[41] J. Thomas, R. Bashyal, S. Goldstein, and E. Suma, "MuVR: A multi-user virtual reality platform," in *IEEE VR*, 2014.

[42] E. Bastug, M. Bennis, M. Médard, and M. Debbah, "Towards interconnected virtual reality: Opportunities, challenges and enablers," *CoRR*, vol. abs/1611.05356, 2016. [Online]. Available: http://arxiv.org/abs/1611.05356
[43] D. Tsoupikova, K. Triandafilou, S. Solanki, A. Barry, F. Preuss, and D. Kamper, "Real-time diagnostic data in multi-user virtual reality post-stroke therapy," in *ACM SIGGRAPH ASIA*, 2016.
[44] "The effects of 3D multi-user virtual environments on freshmen university students' conceptual and spatial learning and presence in departmental orientation," *Computers & Education*, vol. 94, pp. 228 – 240, 2016.
[45] "Learning to work together: Designing a multi-user virtual reality game for social collaboration and perspective-taking for children with autism," *International Journal of Child-Computer Interaction*, vol. 6, pp. 28 – 38, 2015.
[46] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE J. Sel. Areas Commun*, vol. 20, no. 4, pp. 756–767, 2002.
[47] D. Xu, M. Chiang, and J. Rexford, "Link-state routing with hop-by-hop forwarding can achieve optimal traffic engineering," *IEEE/ACM Trans. Netw.*, vol. 19, 2011.
[48] J. W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang, "Joint VM placement and routing for data center traffic engineering," in *IEEE INFOCOM*, 2012.
[49] X. Li and M. J. Freedman, "Scaling IP multicast on datacenter topologies," in *ACM CoNEXT*, 2013.
[50] "Fast reroute from single link and single node failures for IP multicast," *Computer Networks*, vol. 82, pp. 20 – 33, 2015.
[51] S. Brandt, K. T. Frster, and R. Wattenhofer, "On consistent migration of flows in SDNs," in *IEEE INFOCOM*, 2016.
[52] A. Kabbani, B. Vamanan, J. Hasan, and F. Duchene, "Flowbender: Flow-level adaptive routing for improved latency and throughput in datacenter networks," in *ACM CoNEXT*, 2014.
[53] G. Rtvri, J. Tapolcai, G. Enyedi, and A. Csszr, "IP fast reroute: Loop free alternates revisited," in *IEEE INFOCOM*, 2011.
[54] H. Wang, Y. Li, Y. Zhang, and D. Jin, "Virtual machine migration planning in software-defined networks," in *IEEE INFOCOM*, 2015.
[55] V. Shrivastava, P. Zerfos, K. w. Lee, H. Jamjoom, Y. H. Liu, and S. Banerjee, "Application-aware virtual machine migration in data centers," in *IEEE INFOCOM*, 2011.
[56] J. Liu, L. Su, Y. Jin, Y. Li, D. Jin, and L. Zeng, "Optimal VM migration planning for data centers," in *IEEE GLOBECOM*, 2014.
[57] "Virtualknotter: Online virtual machine shuffling for congestion resolving in virtualized datacenter," *Computer Networks*, vol. 67, pp. 141 – 153, 2014.
[58] IBM cloud: SoftLayer. [Online]. Available: http://www.softlayer.com/Data-Centers
[59] Cogent communication. [Online]. Available: http://cogentco.com/en/network/network-map
[60] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree-based vs. structural," in *ACM SIGCOMM*, 2002.
[61] IBM ILOG CPLEX. [Online]. Available: http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/
[62] Y. Li, L. T. X. Phan, and B. T. Loo, "Network functions virtualization with soft real-time guarantees," in *IEEE INFOCOM*, 2016.
[63] Emulab. [Online]. Available: https://www.emulab.net/

APPENDIX A
THE PROOF OF THEOREM 1

Proof: We prove it by a polynomial-time reduction from a variant of the Steiner tree problem, where all the edge costs are positive and satisfy triangular inequality, to the SOF problem. Given any instance G of the Steiner tree problem, we construct a corresponding instance G' of the SOF problem as follows. We first replicate G into G', set $|\mathcal{C}| = 1$ in G', and add one source s into G'. We let root r as the only VM in G' and nodes in U of G as the destinations in D of G'. Root r is connected to s with an edge whose cost is set to an arbitrary value $w > 0$ so as to obtain the instance G'. In the following, we prove $OPT_{G'} = OPT_G + w$. Because the edge $e_{r,s}$ only exists in G' and the solution of G' must contain a subgraph in G, which is also a tree rooted at r and spans all the nodes in U, $OPT_G \leq OPT_{G'} - w$ holds. In addition, $OPT_G \geq OPT_{G'} - w$; otherwise, such a Steiner tree of G plus the edge $e_{r,s}$ becomes a solution of G' with a smaller cost than $OPT_{G'}$. Hence, given OPT_G (or $OPT_{G'}$), we can obtain $OPT_{G'}$ (or OPT_G) by adding (or removing) the edge $e_{r,s}$. The theorem follows. ■

APPENDIX B
THE PROOF OF LEMMA 1

Proof: Consider any three nodes a, b, and c in $\mathcal{G}$. Since $\mathcal{G}$ is a complete graph, the edge between a and b, the edge between b and c, and the edge between a and c form a triangle in $\mathcal{G}$. Clearly, the total cost of the edge between a and b and the edge between b and c must be greater than the cost of the edge between a and c; otherwise, the connection cost of the shortest path between a and c in G must be greater than the total connection cost of the shortest path between a and b and the shortest path between b and c, which is a contradiction. The lemma follows. ■

APPENDIX C
THE PROOF OF THEOREM 2

Proof: There are two possible cases for the last VM u. For the first case, u is one of the last VMs in $\mathcal{F}^{OPT}$, and u is not in the second case. For the first case, let $\mathcal{F}^{OPT}(u)$ be a subgraph of $\mathcal{F}^{OPT}$ that connects u to all the destinations in D. Thus, $c(\mathcal{F}^{OPT}(u)) \leq c(\mathcal{F}_E^{OPT})$. It is worthy to note that the subgraph $\mathcal{F}^{OPT}(u)$ may also span s (and even the other last VMs in $\mathcal{F}^{OPT}$) so as to connect u and all the destinations. Hence, the cost of the minimum Steiner tree that spans u and all the destinations in D must be no greater than $c(\mathcal{F}^{OPT}(u)) \leq c(\mathcal{F}_E^{OPT})$ when u also runs the last VNF in $\mathcal{F}^{OPT}$. On the other hand, the k-stroll problem is NP-Hard and has a 2-approximation algorithm in metric graphs, which satisfy triangular inequality. According to Lemma 1, $W_{\mathcal{G}}$ follows triangular inequality, and the cost of $W_{\mathcal{G}}$ is thereby no greater than twice of the cost of the shortest walk that visits at least $|\mathcal{C}|+1$ distinct nodes, $W_{\mathcal{G}}^{OPT}$, from s to u in $\mathcal{G}$. Since the cost of $W_{\mathcal{G}}^{OPT}$ is equal to the minimum setup and connection costs of the walk that visits at least $|\mathcal{C}|$ VMs from s to u in G, the cost of $W_{\mathcal{G}}^{OPT}$ is bounded by the total setup and connection costs of the walk from s to u in $\mathcal{F}^{OPT}$. Therefore, the cost of $W_{\mathcal{G}}$ is bounded by $2c(\mathcal{F}^{OPT})$.

On the other hand, the connection cost of connecting u to all destinations is bounded by $\rho_{ST} \cdot c(\mathcal{F}_E^{OPT})$. Thus, the cost of the service overlay forest with the last VM u is bounded by $(2+\rho_{ST})c(\mathcal{F}^{OPT})$ as u runs $f_{|\mathcal{C}|}$ in $\mathcal{F}^{OPT}$. Note that SOFDA-SS constructs a service overlay forest for every possible last VM u and chooses the forest with the minimum cost. Since at least one VM runs $f_{|\mathcal{C}|}$ in $\mathcal{F}^{OPT}$, the cost of the service overlay forest generated by SOFDA-SS is bounded by $(2 + \rho_{ST})c(\mathcal{F}^{OPT})$. The theorem follows. ■

APPENDIX D
SCENARIOS WITH SETUP COSTS ON SOURCES

In the following, we extend SOFDA-SS to support the case with a setup cost assigned to the source. Let $c(s)$ denote the cost to enable source s. Let the cost of the edge between nodes v_1 and v_2 in $\mathcal{E}$

$$c(v_1, v_2) = \sum_{(a,b)\in P} c((a,b)) + \begin{cases} c(s)+c(u) & \text{if } v_1 = s, v_2 = u \text{ or } v_1 = u, v_2 = s \\ \frac{c(s)+c(u)+c(v_2)}{2} & \text{else if } v_1 = s, v_2 \neq u \text{ or } v_1 = u, v_2 \neq s \\ \frac{c(v_1)+c(s)+c(u)}{2} & \text{else if } v_1 \neq s, v_2 = u \text{ or } v_1 \neq u, v_2 = s \\ \frac{c(v_1)+c(v_2)}{2} & \text{otherwise,} \end{cases}$$

where u and P denote the last VM and the shortest path between nodes v_1 and v_2 in G, respectively. Due to a similar reason in Section IV, we set the cost of the edges in $\mathcal{E}$ in a similar way but consider the source cost $c(s)$. The approximation ratio also holds, and the proof is similar to Theorem 2.

APPENDIX E
PSEUDO CODES

Procedure 1 Instance Construction of the k-Stroll Problem

Input: A network $G = (V, E)$, a source $s \in V$, a set of VMs $M \subseteq V$, and a (last) VM $u \in M$

$\mathcal{V} \leftarrow M \cup \{s\}$;
$\mathcal{E} \leftarrow \{(x, y)|x, y \in \mathcal{V}\}$;
for each edge e in $\mathcal{E}$ between nodes v_1 and v_2 **do**
  $P \leftarrow$ the shortest path between v_1 and v_2;

$$c(v_1, v_2) \leftarrow \sum_{(a,b)\in P} c((a,b)) + \begin{cases} \frac{c(u)+c(v_2)}{2} & \text{if } v_1 = s, \\ \frac{c(v_1)+c(u)}{2} & \text{else if } v_2 = s, \\ \frac{c(v_1)+c(v_2)}{2} & \text{otherwise;} \end{cases}$$

end for
return $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$;

Procedure 2 Identification of the Walk with $|\mathcal{C}|$ VMs

Input: A network $G = (V, E)$, a source $s \in V$, a set of VMs $M \subseteq V$, a (last) VM $u \in M$, and the length of the chain of VNFs $|\mathcal{C}|$

construct an instance $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of the k-stroll problem using Procedure 1 with input parameters G, s, M, and u;
obtain a walk $W_{\mathcal{G}}$ that visits at least $|\mathcal{C}| + 1$ distinct nodes from s to u in $\mathcal{G}$ using the 2-approximation algorithm for the metric version of the k-stroll problem in [29];
obtain a walk $W'_{\mathcal{G}} = (u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$, which visits exactly $|\mathcal{C}| + 1$ distinct nodes from s (=u_1) to u (=$u_{|\mathcal{C}|+1}$), in $\mathcal{G}$ from $W_{\mathcal{G}}$ by repeatedly removing a node until $|\mathcal{C}| + 1$ nodes remain;
obtain the walk $W_G = (u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$ from s (=u_1) to u (=$u_{|\mathcal{C}|+1}$) by combining the shortest paths between u_j and u_{j+1} for $j = 1, 2, \cdots, |\mathcal{C}|$ in G;
return $W_G = (u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$;

Algorithm 1 $(2+\rho_{ST})$-Approximation Algorithm for the SOF Problem with Single Source

Input: A network $G = \{V, E\}$, a set of destinations $D \subseteq V$, a source $s \in V$, a set of VMs $M \subseteq V$, and a chain of VNFs $\mathcal{C} = (f_1, f_2, \cdots, f_{|\mathcal{C}|})$

$C \leftarrow \infty$;
for each VM u **do**
  obtain the walk $W_G = (u_1, u_2, \cdots, u_{|\mathcal{C}|+1})$ from s (=u_1) to u (=$u_{|\mathcal{C}|+1}$) using Procedure 2 with input parameters G, s, M, u, and $|\mathcal{C}|$;
  construct F_{tmp} with cost C_{tmp} by connecting s with u by W_G, running $f_1, f_2, \cdots, f_{|\mathcal{C}|}$ at $u_2, \cdots, u_{|\mathcal{C}|+1}$, respectively, and connecting u with all destinations using the approximation algorithm for the Steiner Tree problem in [20];
  if $C_{tmp} < C$ **then**
    $F \leftarrow F_{tmp}$;
    $C \leftarrow C_{tmp}$;
  end if
end for
return F;

Procedure 3 Instance Construction of the Steiner Tree Problem

Input: A network $G = (V, E)$, a set of sources $S \subseteq V$, a set of VMs $M \subseteq V$, and a chain of VNFs $\mathcal{C} = (f_1, f_2, \cdots, f_{|\mathcal{C}|})$

$\mathbb{V} \leftarrow V \cup \{\hat{s}\} \cup \mathbb{V}_{\mathbb{S}} \cup \mathbb{V}_{\mathbb{M}}$, where $\mathbb{V}_{\mathbb{S}}$ is a set of the duplicate $\hat{v}$ of each $v \in S$, $\mathbb{V}_{\mathbb{M}}$ is a set of the duplicate $\hat{v}$ of each $v \in M$, and $\hat{s}$ is a virtual source;
$\mathbb{E} \leftarrow E \cup \mathbb{E}_{\hat{s}\mathbb{S}} \cup \mathbb{E}_{\mathbb{S}\mathbb{M}} \cup \mathbb{E}_{\mathbb{M}M}$, where $\mathbb{E}_{\hat{s}\mathbb{S}}$ is a set of the edges $(\hat{s}, \hat{v}$ for all $\hat{v} \in \mathbb{V}_{\mathbb{S}}$, $\mathbb{E}_{\mathbb{S}\mathbb{M}}$ is a set of the edges $(\hat{v}, \hat{u})$ for all $\hat{v} \in \mathbb{V}_{\mathbb{S}}$ and $\hat{u} \in \mathbb{V}_{\mathbb{M}}$, and $\mathbb{E}_{\mathbb{M}M}$ is a set of the edges $(v, \hat{v})$ for all $v \in M$;
for each edge e in $\mathbb{E}_{\hat{s}\mathbb{S}} \cup \mathbb{E}_{\mathbb{M}M}$ **do**
  the weight of $e \leftarrow 0$;
end for
for each edge $(\hat{v}, \hat{u})$, where $\hat{v} \in \mathbb{S}$ and $\hat{u} \in \mathbb{M}$, in $\mathbb{E}_{\mathbb{S}\mathbb{M}}$ **do**
  the weight of $(\hat{v}, \hat{u}) \leftarrow$ the cost of the walk that visits $|C|$ VMs from $v \in S$ to $u \in M$ in G using Procedure 2 with input parameters G, v, M, u, and $|C|$;
end for
return $\mathbb{G} = \{\mathbb{V}, \mathbb{E}\}$;

Algorithm 2 $3\rho_{ST}$-Approximation Algorithm for the SOF problem

Input: A network $G = \{V, E\}$, a set of destinations $D \subseteq V$, a set of sources $S \subseteq V$, a set of VMs $M \subseteq V$, and a chain of VNFs $\mathcal{C} = (f_1, f_2, \cdots, f_{|\mathcal{C}|})$

construct an instance $\mathbb{G} = \{\mathbb{V}, \mathbb{E}\}$ of the Steiner Tree problem using Procedure 3 with input parameters G, S, M, and $\mathcal{C}$;
obtain a Steiner tree $\mathbb{T}$ in $\mathbb{G}$ using the ρ_{ST}-approximation algorithm for the Steiner Tree problem in [20];
$F \leftarrow \emptyset$;
for each pair of $v \in S$ and $u \in M$ in G such that $\hat{v}$ (duplicate of v) is the parent of $\hat{u}$ (duplicate of u) in tree $\mathbb{T}$ with root $\hat{s}$ **do**
  obtain the walk W that visits $|C|$ VMs between v and u in G using Procedure 2 with input parameters G, v, M, u, and $|C|$;
  add W to F using Procedure 4 with input parameters F, W, and $\mathcal{C}$;
end for
add all VMs, switches, and links in $\mathbb{T} \cap G$ to F;
return F;

Procedure 4 Augment of Service Overlay Forest

Input: A service overlay forest F, a walk W_G in G, and a chain of VNFs $\mathcal{C} = (f_1, f_2, \cdots, f_{|\mathcal{C}|})$

obtain $W = (u_1, u_2, \cdots, u_n)$ from $W_G = (v_1, v_2, \cdots, v_n)$, where u_i:
- is the (clone of) VM v_i running a VNF in F if v_i is a VM in G and there is already (one clone of) VM v_i running a VNF in F and
- is the newly-added (clone of) VM (or a switch) v_i in F otherwise;

augment F with W;
if W experiences the *VNF conflict* with at least one walk in F **then**
  for each walk W_k that experiences the *VNF conflict* with W by backtracking W **do**
    if W (bewteen source s and VM v) experiences the *VNF conflict* with W_k (bewteen source s_1 and VM v_1) at VM u for the first time and the sequence number of the VNF (f_j) on W is not greater than that of the VNF (f_i) on W_k at VM u **then**
      attach W to W_k through u by changing W to the concatenation of the sub-walk of W_k from s_1 to u (on which $f_1, f_2, \cdots, f_i$ are run in sequence) and the sub-walk of W from u to v (on which $f_{i+1}, f_{i+2}, \cdots, f_{|C|}$ are run in sequence);
      return F;
    else if there exists another VM w such that W experiences the *VNF conflict* with W_k at w and the sequence number of the VNF (f_h) at w on W_k is not smaller than that of the VNF (f_j) at u on W **then**
      attach W to W_k through w by changing W to the concatenation of the sub-walk of W_k from s_1 to w (on which $f_1, f_2, \cdots, f_h$ are run in sequence), the sub-walk of W from w to u, and the sub-walk of W from u to v (on which $f_{h+1}, f_{h+2}, \cdots, f_{|C|}$ are run in sequence);
      return F;
    else
      attach W_k to W through u by changing W_k to the concatenation of the sub-walk of W from s to u (on which $f_1, f_2, \cdots, f_j$ are run in sequence) and the sub-walk of W_k from u to v_1 (on which $f_{j+1}, f_{j+2}, \cdots, f_{|C|}$ are run in sequence);
    end if
  end for
end if
return F;

TABLE III
TABLE OF NOTATIONS

G	the input network
V	the set of nodes in the network
M	the set of available VMs in the network
U	the set of switches in the network
S	the set of sources in the network
D	the set of destinations in the network
$c(\cdot)$	the cost function
$\mathcal{C}$	the demanded service chain
f_i	the i^{th} VNF in chain $\mathcal{C}$
$\mathcal{F}^{OPT}$	the optimal service forest
$\mathcal{F}_E^{OPT}$	the set of used edges in the optimal forest
$\mathcal{F}_M^{OPT}$	the set of used VMs in the optimal forest
$\mathcal{F}^{OPT}(u)$	the subgraph of $\mathcal{F}^{OPT}$ which connects node u and all the destinations
D_v^{OPT}	the set of destinations whose source is v in $\mathcal{F}^{OPT}$
m_v^{OPT}	the representative last VM among all last VM on the walks from source v in $\mathcal{F}^{OPT}$
T_v	the Steiner tree that spans m_v^{OPT} and the destinations in D_v^{OPT} for source v in G
$\mathcal{G}$	the constructed instance of k-stroll problem
$\mathcal{V}$	the nodes in $\mathcal{G}$
$\mathcal{E}$	the edges in $\mathcal{G}$
$\mathcal{A}$	the 2-approximation for k-stroll problem
$W_{\mathcal{G}}$	the output walk in $\mathcal{G}$ by $\mathcal{A}$
$W'_{\mathcal{G}}$	the walk that visits exactly $\|\mathcal{C}\|$ distinct nodes in $\mathcal{G}$ obtained from $W_{\mathcal{G}}$
W_G	the walk in G that visits $\|\mathcal{C}\|$ VMs in G
$\mathbb{G}$	the constructed instance of Steiner tree problem
$\mathbb{V}$	the set of nodes in $\mathbb{G}$
$\mathbb{E}$	the set of edges in $\mathbb{G}$
$\hat{s}$	the virtual source in $\mathbb{G}$